The Caltech Fish Counting Dataset: A Benchmark for Multiple-Object Tracking and Counting

Justin Kay^{1,5}, Peter Kulits¹, Suzanne Stathatos¹, Siqi Deng², Erik Young³, Sara Beery¹, Grant Van Horn⁴, and Pietro Perona^{1,2}

 1 California Institute of Technology $\ ^2$ AWS AI Labs $\ ^3$ Trout Unlimited 4 Cornell University $\ ^5$ Ai.Fish

Abstract. We present the Caltech Fish Counting Dataset (CFC), a large-scale dataset for detecting, tracking, and counting fish in sonar videos. We identify sonar videos as a rich source of data for advancing low signal-to-noise computer vision applications and tackling domain generalization in multiple-object tracking (MOT) and counting. In comparison to existing MOT and counting datasets, which are largely restricted to videos of people and vehicles in cities, CFC is sourced from a natural-world domain where targets are not easily resolvable and appearance features cannot be easily leveraged for target re-identification. With over half a million annotations in over 1,500 videos sourced from seven different sonar cameras, CFC allows researchers to train MOT and counting algorithms and evaluate generalization performance at unseen test locations. We perform extensive baseline experiments and identify key challenges and opportunities for advancing the state of the art in generalization in MOT and counting.

Keywords: Detection, Tracking, Counting, Video Dataset

1 Introduction

Diverse and high-quality datasets collected from the natural world have enabled progress on fundamental computer vision tasks such as fine-grained visual categorization [9, 10, 11, 67, 82, 102, 107, 108, 109] and individual re-identification [50, 69, 86]. This progress has had valuable impact in the same natural-world domains, and automated visual systems are now used in the field every day by ecologists, citizen scientists, and conservationists to improve the accuracy and efficiency of biodiversity monitoring efforts around the globe [1, 12, 41, 66, 106, 107]. However, the range of computer vision tasks that these systems can perform is still limited, with most current algorithms and their supporting datasets focusing largely on visual classification in relatively high-quality imagery. Methods developed for existing datasets can fail to transfer to tasks where video analysis or non-RGB imagery is involved.

We present the Caltech Fish Counting Dataset (CFC), a large video dataset containing over half a million annotations for detecting, tracking, and counting

migrating fish in sonar video. In addition to providing a challenging benchmark in a novel application domain, the dataset allows for detailed study in three areas that have received limited attention from the computer vision community and lack supporting benchmarks:

1. Multiple-object tracking (MOT) in natural environments with animal targets. Most existing MOT datasets focus on human [34, 68, 78] and vehicle [44, 101, 113, 120] tracking in cities. Animal targets provide a rich source of variability for developing trackers that are not biased toward urban domains, in addition to having numerous beneficial applications in conservation [106], neuroscience [76], and animal husbandry [42]. Furthermore, many state-of-the-art methods make extensive use of visual re-identification for performing association [88, 116, 126]. In contrast, CFC provides a benchmark sourced from low signal-to-noise recording equipment in challenging natural-world environments where tracking targets are difficult to resolve from background clutter and each other, making frame-to-frame visual association less effective.

The dataset is large, well-annotated, and challenging. It consists of 1,567 video sequences sourced from seven different sonar cameras on three rivers located in the U.S. states of Alaska and Washington. The videos are single-channel (i.e. grayscale), vary in resolution from 288x624 to 1,086x2,125, have frame rates between 6.7 and 13.3 frames per second, and are an average of 336 frames (38 seconds) in duration. Tracking annotations were collected through a paid annotation service for 8,254 fish across 527k frames, totaling 516k bounding boxes in 16.7 hours of video.

2. Video-based counting. Existing video counting benchmarks focus on crowd counting in urban environments [25, 26, 40, 40, 114, 115] and emphasize density estimation over trajectory-based counting. Methods developed for these datasets have limitations in applications where information about individuals, such as size or direction of motion, is required, since crowd density is largely treated as a regional feature [95]. The community is in need of a benchmark which can support both tracking and counting concurrently. CFC provides both a challenging MOT benchmark and an evaluation protocol for video-based counting that is motivated by a real-world metric. Ground-truth detections, trajectories, and counts are provided for every video sequence.

3. Generalization of tracking and counting methods to new domains. While generalization in computer vision has been extensively studied in object recognition [10,30,63,131], it is still relatively understudied within MOT and counting. CFC presents significant generalization challenges that are highlighted by the dataset design, providing a strong benchmark for the study of generalization and efficient adaptation in the context of MOT and counting. We enable this study by constraining training data to a single camera location, while sourcing test data from a variety of different out-of-sample rivers and cameras.

Finally, CFC is the first annotated video dataset sourced from the domain of fish counting in sonar, an application area with significant impacts in conservation ecology. Salmon are keystone species that support at least 137 other animal species and provide food and nutrients to a wide range of ecosystems during their

Table 1: Comparison of video tracking and counting datasets. CFC is the first dataset that supports all three tasks of interest: detection, tracking, and counting, with more tracking annotations than existing animal-centric tracking datasets. (* indicates annotations are points, not bounding boxes)

Dataset	Vids	Frames	Annos	Animals	Detect	Track	Count
3D-ZeF [87]	8	28,800	86,400	\checkmark (100%)	✓	\checkmark	
BIRDSAI [19]	48	62,400	154,000	√ (78%)	\checkmark	\checkmark	
GMOT-40 [7]	40	9,643	256,341	√ (38%)		\checkmark	
MOT16 [78]	14	11,235	292,733		\checkmark	\checkmark	
MOT20 [34]	8	$13,\!410$	$1.65 \mathrm{M}$		\checkmark	\checkmark	
UA-DETRAC [113]	100	140,000	$1.21 \mathrm{M}$		\checkmark	\checkmark	
TAO [33]	$2,\!907$	4.44M	332,401	√ (~10%)	\checkmark	\checkmark	
AnimalDrone [133]	162	$53,\!644$	$4.05 M^{*}$	√ (100%)			\checkmark
DroneCrowd [115]	112	$33,\!600$	$4.86 M^{*}$			\checkmark	\checkmark
Crossing-line [129]	5	3,100	5,900*			\checkmark	\checkmark
FDST [40]	100	$15,\!000$	394,081*				\checkmark
Iowa DOT [81]	200	90,000	0				\checkmark
Mall [27]	1	2,000	$62,315^*$				\checkmark
UCSD [25]	1	2,000	$49,885^{*}$				\checkmark
WorldExpo [121]	$1,\!132$	3,980	199,923*				✓
CFC	1,567	527,215	515,933	\checkmark (100%)	\checkmark	\checkmark	\checkmark

seasonal migration [43]. Sonar imaging provides a non-invasive way to monitor *escapement*—the number of salmon returning home each season to spawn—helping inform sustainable fisheries management. Automation using computer vision could enable current sonar-based monitoring programs to scale from a few locations to entire watersheds. We hope that our dataset will encourage computer vision researchers to work on this high-impact challenge.

Our contributions are: (1) a large and challenging dataset for tracking and counting in video that enables the study of generalizing algorithms to new locations; (2) an evaluation protocol that mimics the procedure used by field technicians when manually counting fish in sonar video; (3) a baseline method that utilizes a novel input structure to improve generalization performance at unseen test locations. The dataset and evaluation code are available here.

2 Related Work

Multiple-Object Tracking (MOT). Several popular benchmarks have supported recent progress in MOT, particularly in the domains of pedestrian and vehicle tracking [34, 44, 68, 78, 101, 113, 120]. Large-scale benchmarks for tracking animals in the wild are less common. TAO [33] and GMOT-20 [7] focus on tracking all foreground objects with limited class information, and include some animal tracking sequences. The BIRDSAI [19] dataset focuses on human

and animal tracking in the wild, and—similar to ours—contains non-RGB (in their case, thermal infrared) sequences, though it is sourced from moving aerial cameras rather than in-situ monitoring devices. With over half a million MOT annotations, the CFC dataset is larger than existing benchmarks for animal tracking in the wild, with a unique additional focus on video-based counting and generalization challenges. See Tab. 1 for a comparison with prior work.

Recent advancements in object detection [54, 134] have helped popularize the tracking-by-detection paradigm in MOT, which divides the tracking problem into two steps: (1) an object detector predicts object locations (e.g. bounding boxes) in each frame, and (2) a tracker associates detections over time into object trajectories. While there has been significant progress in recent years (see recent surveys [29,72]), there are two notable shortcomings. First, the trackingby-detection paradigm implicitly assumes that detection is possible and accurate in each frame. This is not universally valid, and our dataset can stimulate research in algorithms that do not rely on this assumption due to the difficulty in resolving fish locations frame-to-frame. Second, much recent progress can be attributed to the development of complex visual-feature representations for target re-identification. These techniques are often domain-specific according to the benchmark dataset. For example, [130] use a generative model to create synthetic pedestrian data consisting of various combinations of person appearance and structure information to achieve state-of-the-art performance, a technique that has been adopted by other top-performing MOT methods [52]. Our dataset introduces a challenging MOT benchmark in which individuals are visually indistinct, offering little opportunity to make use of visual features for individual re-identification, which can motivate the development of tracking methods that are not dependent upon complex appearance matching.

Counting in Video. Datasets for object counting are predominantly imagebased [4, 8, 22, 45, 53, 55, 57, 60, 83, 89, 97, 110, 112, 117, 121, 122, 127]. Video-based counting datasets are more limited. Most are focused on estimating the number of people [25, 26, 40, 40, 114, 115] or animals [133] in crowded scenes, combining the challenges of crowd-density estimation and camera motion compensation. While a limited number of video counting datasets incorporate trajectory information [115, 129], existing benchmarks primarily model object locations as points and do not contain bounding box labels. CFC supports the study of all three tasks (detection, tracking, and counting), while containing over five times the number of annotated video frames as existing video counting benchmarks (see Tab. 1).

Methods for video-based counting can be roughly divided into *regression-based* methods, *density-based* methods, and *detection-based* methods. Regression-based methods attempt to predict counts directly by mapping image features to counting numbers [59, 79], while density-based methods predict per-pixel crowd density in each frame and then analyze densities over time to obtain counts [4,84,123,127,129]. These methods are typically designed for counting large numbers of densely clustered objects where individual object detection is challenging. In contrast, detection-based methods utilize object detection in each frame to localize objects of interest and count them over time. These approaches typically employ a tracking-by-detection pipeline followed by counting based on either *region-of-interest* (ROI) [24, 74] or *line-of-interest* (LOI) [62, 73, 129]. ROI-based methods attempt to estimate the number of objects passing through a subregion of the frame, such as a traffic lane or onramp in the case of vehicle counting [24]. LOI-based methods instead draw a virtual line through the field of view, counting objects when their trajectories intersect this line. In this work we use LOI-based counting, as it matches the approach currently used by fishery managers for counting fish in sonar [61].

Animal-Centric Datasets. Existing computer vision datasets in animal ecology primarily target the tasks of species classification [9,11,82,102,107,108,109], detection [2, 6, 9, 10, 20, 31, 35, 57, 85, 90, 94, 99, 102, 103, 104, 119, 128], and individual identification [50, 69, 86]. These datasets consist predominantly of RGB imagery where visual features are key signals for recognition; in contrast, our dataset consists of single-channel sonar video, in which the animals of interest are difficult to distinguish from background, debris, and each other.

For this reason CFC shares some characteristics with video datasets for studying animal behavior, including mice [3,5,17,46,47,51,100], rats [32], flies [37,38], bees [16,21,75,92], and fish [23,87,132]. As in our dataset, the visual similarities between individuals in these datasets mean that tracking must rely heavily on prediction of motion or behavior rather than using visual features for association. Our work, however, entails additional challenges not typically present in laboratory tracking and behavior study, such as complex background, difficult frame-by-frame detection, and unknown numbers of individuals in a scene.

Generalization in Computer Vision. Domain generalization is a type of domain shift—i.e. a difference in training and test data distributions—in which training and test data come from distinct, but related, domains. For example, in our dataset the domain generalization challenge comes from training and test data sourced from distinct sonar camera deployments on different rivers. Crucially, as opposed to *domain adaptation*, in which data from the test domain is available during model training, in domain generalization data from the test domain is considered inaccessible [15, 80, 105]. Within computer vision this has been most extensively studied for the task of object recognition [30,48,118,131], with a number of supporting datasets [10, 39, 63, 93]. While there has been some study of domain generalization in other computer vision tasks such as semantic segmentation [28,49,124] and action recognition [65,98,111], it has been relatively understudied in the context of MOT and counting. Some MOT and counting datasets do represent domain generalization challenges in their test sets—for example, in MOT20 [34] one of the three test locations is from a new location, and in UA-DETRAC [113] (MOT) and WorldExpo [121] (crowd counting) all test data is from different locations than training data. Our dataset makes it possible to evaluate generalization for both MOT and video-based counting concurrently, while providing more out-of-distribution test videos than existing options.

Imaging Sonar. Only a very limited amount of annotated sonar imaging data has been released. [77] collected 524 sonar video clips with video-level species labels, [96] collected a small dataset of 143 sonar images to test image-level





Fig. 1: Illustration of camera strata and data split. (A) Three example frames, one from each range window ("stratum") of the far-range camera at the KR location. Sonar cameras typically cycle between multiple strata periodically. White bounding boxes are ground-truth fish locations. (B) All cameras in the dataset. There are seven cameras total distributed among the five locations in the dataset, each with between one and three strata. Data split is indicated by the colored circles, and the number of training, validation, and testing sequences are indicated for each camera/stratum.

classification of fish and dolphin species, and [70] collected dot annotations for counting fish in 537 sonar images. Our work is the first to release detection and tracking annotations for fish in sonar, and it is several orders of magnitude larger than existing sonar datasets.

3 Dataset

Here we describe how we collected, annotated, and split CFC, and give an overview of common challenges.

Data Collection. The dataset was curated from 2,056 hours of sonar video obtained from the Alaska Department of Fish and Game, the U.S. National Marine Fisheries Service, the U.S. National Park Service, and the Lower Elwha Klallam Tribe. It contains video from five distinct locations which we use to study out-of-sample performance: three locations on the Kenai River in Alaska, which we refer to as **KL** (Kenai Left Bank), **KR** (Kenai Right Bank), and **KC** (Kenai Channel); one location, NU, on the Nushagak River in Alaska; and one location, EL, on the Elwha River in Washington (sonar configurations shown in Fig. 1).

6

 $\overline{7}$



Fig. 2: Example frames and common challenges in sonar video. Groundtruth fish locations are boxed in white. (A) Close range, high operating frequency: ideal conditions, fish is large, visible, and well-defined. Some speckle noise is still present. (B) Far range, low operating frequency: The fish are small and very coarsely defined due to scattering of sound waves at long range. (C) Background texture: The riverbed is very visible, occluding fish. (D) Shadows: Fish cast acoustic shadows which may occlude one another. (E) Sediment: Dirt, debris, and glacial silt occlude fish. (F) Target density: Dense crowds of fish, intersecting trajectories and occlusion.

The data had already been analyzed by experts ("manually marked") to obtain fish counts. Most of the video contained no fish. Since our focus is on detection, tracking, and counting, we used the manual markings to extract shorter 200–300 frame video clips known to contain fish. If any of these clips overlapped, we merged them into one longer clip. In total we extracted 1,233 clips from the Kenai River, containing 4,300 fish; 262 clips from the Elwha River, containing 884 fish; and 72 clips from the Nushagak River, containing 3,070 fish.

Annotation. We hired a third-party annotation service to collect multipleobject tracking annotations for all fish in the extracted clips. Annotators were provided with the raw video clips and instructed to box all visible fish tightly using the vatic.js GUI [18]. For any stationary fish, they were required to annotate every fifth frame, and we interpolated between those annotations in the intermediate frames; for all other tracks, all bounding boxes in all frames were annotated manually. The annotation service had their own internal quality-management procedures whereby multiple annotators inspected each clip before it was finalized. In total, 515,933 bounding boxes for 8,254 fish tracks were collected in 527,215 frames from seven different cameras.

Data Split. We designed a dataset splitting protocol that allows us to study generalization to new locations, a known challenge for current computer vision methods [10,64]. Our test data comes from deployment locations never seen during training or validation. We chose KL as our training and validation location due to its sufficient size for model training. Data from this location spans 16

days in total. We selected one of these days at random to hold out as a validation set, and the other 15 days serve as our training set. This gives us 162,680 training images containing 1,762 tracks and 132,220 bounding boxes, and 30,518 validation images containing 207 tracks and 18,565 bounding boxes.

The other locations (KR, KC, NU, and EL) serve as test sets for evaluating generalization performance under different conditions. These locations cover a range of generalization scenarios that may be faced in the real world: KR includes a new camera deployment on the same body of water; KC includes a new deployment in a nearby, but separate, body of water; NU and EL include deployments on new rivers in different geographic regions with different species distributions. The data split is illustrated in Fig. 1. In total the dataset contains 334,017 test images with 6,285 tracks and 365,148 bounding boxes. For all experiments we report results on all test locations individually.

Challenges. We have identified a number of challenges inherent to detecting and tracking fish using sonar, which we have illustrated in Fig. 2. Some of these challenges are constant across all data in this domain (e.g. speckle noise and shadows), while some vary across locations due to hardware settings or environmental factors, presenting generalization challenges (e.g. presence of sediment, riverbed shape and texture, and hardware operating frequency). More details on the causes of these challenges can be found in the supplemental material.

4 Metrics

4.1 Counting Protocol

We follow the counting procedure used by field technicians when counting fish in sonar video [61]. A vertical line-of-interest (LOI) is drawn in the middle of the frame, and a fish is considered to have moved left or right if its trajectory start and end positions are on different sides of the LOI. Note that not every fish in a clip will cross the LOI. Some fish are stationary throughout the entire clip, while others enter and exit on the same side of the frame without crossing the LOI. These fish are excluded from the count totals, which matches the protocol used by the field technicians.

4.2 Counting Metric

In the target application, fish-counting error is measured as the sum of upstream and downstream counting errors, and error is normalized separately at each river to account for variations in fish abundance. We classify direction of movement as "left" or "right" rather than "upstream" or "downstream" to make our system agnostic to the orientation of the camera. Based on this we define the absolute counting error for the *i*th video clip, E_i , as the sum of the absolute left and right (i.e. upstream and downstream) counting errors:

$$E_i = |z_{left_i} - \hat{z}_{left_i}| + |z_{right_i} - \hat{z}_{right_i}| \tag{1}$$

where z_i and \hat{z}_i are the predicted and ground-truth counts for clip *i*, respectively. Overall counting error is reported as normalized Mean Absolute Error (nMAE):

$$nMAE = \frac{\frac{1}{N}\sum_{i=1}^{N}E_i}{\frac{1}{N}\sum_{i=1}^{N}\hat{z}_i} = \frac{\sum_{i=1}^{N}E_i}{\sum_{i=1}^{N}\hat{z}_i}$$
(2)

where N is the number of video clips at a given location and $\hat{z}_i = \hat{z}_{left_i} + \hat{z}_{right_i}$ is the ground-truth count for clip *i*.

We choose a metric that is normalized *per location*, giving *equal weight to each fish* at a given location regardless of variance in fish density across clips at that location. We do this because video clips in CFC are arbitrarily generated and are much shorter in duration than videos in the target application, thus *per-clip* normalization—i.e. giving equal weight to each *clip*—would not be appropriate. We consider this the main metric for the dataset, since it is the most important in the target application. Achieving at most 10% counting error on a river would make an algorithm on par with human experts and feasible for augmenting counting in the field [61].

4.3 Detection and Tracking Metrics

In addition to counting, CFC has been annotated to measure detection and tracking performance as well. For our detection metric we choose the PascalVOC [36] evaluation of mean Average Precision with IoU ≥ 0.5 (AP50). For tracking, we report the CLEAR [13], IDF1 [91], and HOTA [71] metrics.

The CLEAR MOT metrics [13] are computed *per-frame* by matching detections from predicted tracks with ground-truth detections. This matching allows the number of true positive (TP), false positive (FP), and false negative (FN) detections to be computed using an IoU threshold between predicted and ground-truth boxes. Recall (CLR_Re) and precision (CLR_Pr) are defined as normal according to the number of TPs, FPs, and FNs, and Multiple Object Tracking Accuracy (MOTA) is defined as:

$$MOTA = \frac{TP - FP - IDSW}{TP + FN} \tag{3}$$

Where IDSW is the sum of all "ID switches", i.e. the number of times a predicted track changes its matched ground-truth track and vice versa. In practice, $IDSW \ll TP$ and MOTA becomes predominantly a measure of detection quality in the tracks. Since this scoring occurs per-frame, MOTA does not capture long-term tracking performance.

In contrast, IDF1 [91] first computes a global (per-video-clip) track matching, i.e. a bipartite matching between all ground-truth and predicted tracks. From this matching, ID true positives (IDTP), false positives (IDFP), and false negatives (IDFN) are computed, and the IDF1 score is defined as the harmonic mean of ID Precision (IDP) and ID Recall (IDR):

$$IDF1 = 2 \cdot \frac{IDP \times IDR}{IDP + IDR}, \ IDP = \frac{IDTP}{IDTP + IDFP}, \ IDR = \frac{IDTP}{IDTP + IDFN}$$
(4)





Fig. 3: Baseline (Sec. 5.1) and Baseline++ (Sec. 5.3) results on CFC. All results are displayed in terms of error, i.e. the lower-left of each plot represents the best performance. Baseline++ improves performance on all tasks across all locations and reduces generalization gaps between the validation location and testing locations. Generalization challenges remain, most notably at KC and EL.

Since this matching is restricted to be static for the length of the clip, IDF1 is a measure of long-term tracking performance.

HOTA [71] is designed to measure both short-term and long-term tracking performance. It is the geometric mean of a detection score (DetA) and an association score (AssA), each defined as the Jaccard index of detection/association TPs, FPs, and FNs. Thus it can be easily decomposed into detection/association precision and recall (DetRe and DetPr, and AssRe and AssPr, respectively). In this formulation, AssRe is inversely correlated with the number of track splits, while AssPr is inversely correlated with the number of track merges.

5 Experiments

We evaluate state-of-the-art methods on CFC to provide a baseline for future work and give insight into the generalization challenges of object detection, multiple-object tracking, and counting. In Sec. 5.1 we propose a tracking-by-detection approach to fish counting which allows us to evaluate each of these tasks, and study its performance. In Sec. 5.2 we perform ablation studies and investigate the upper bounds of this approach and its generalization capabilities. In Sec. 5.3 we introduce an improved baseline method to address these challenges, establish the state of the art on CFC, and discuss remaining challenges.

5.1 Baseline

Our baseline method uses the YOLOv5 [56] object detector and SORT [14] tracker. Trajectories are then analyzed as described in Sec. 4.1 to predict counts.

We chose YOLOv5 after an initial architecture search. These experiments as well as training settings are included in the supplemental material. We chose SORT because it has shown to be a popular and robust tracker across a range of applications, recently achieving state-of-the-art performance on several MOT datasets with minor modifications [125]. It performs tracking using a motion



Fig. 4: Baseline and Baseline++ error analysis and comparison. (A)–(B) Example frames at the KC and EL locations. Note the large reduction in FP and FN detections between the two methods. (C)–(H) Breakdowns of MOTA, IDF1, HOTA, and nMAE into component submetrics. Arrows point from Baseline results to Baseline++ results; Baseline++ markers are larger and have a black edge. All results are displayed in terms of error. This breakdown shows that the low tracking scores of our baseline method at out-of-distribution locations are predominantly due to FP detections (i.e. low detection precision) that cause low CLR_Pr (C), low IDP (D), and low DetPr leading to low DetA (F)–(G). The proposed Baseline++ method succesfully targets a large portion of these FPs, improving all tracking metrics and submetrics. (E) nMAE decomposed into undercounting and overcounting errors, normalized by ground-truth counts. Baseline++ significantly reduces both types of errors.

model based on the Kalman filter [58] without using appearance information for association. We verified our hypothesis that appearance features are not a strong signal for association in CFC by training a visual re-identification model that performed poorly on our validation set. More details of this experiment are provided in the supplemental material. See Figs. 3 and 4 for our performance on the various locations with this baseline method. A tabular version of these results is included in the supplemental material.

Analysis. Our baseline method performs on par with human experts at the location where it is trained, with a validation counting error of less than 5%. Counting generalization performance is best at KR, which matches intuition given that this data is sourced from a nearby location to the training set (KL). However, generalization performance at the other locations is quite poor. We examine the two most challenging locations, KC (53% nMAE) and EL (32.3% nMAE), in Fig. 4. At KC, false positive (FP) and false negative (FN) detections are caused by the presence of complex background information (Fig. 4A), while



12

J. Kay et al.

Fig. 5: Baseline and upper bound tracking and counting performance. Baseline methods are compared to methods that use various amounts of ground-truth data. X axis key: Baseline, Perfect Tracker, Baseline++, Perfect Tracker++, Perfect Detector, Oracle. Note the performance difference between Baseline and Perfect Detector: there is a large performance boost when using ground-truth detections. Using a tracker that can perfectly generalize to different locations improves performance as well, but not as much as the Perfect Detector.

at EL the errors are overwhelmingly FPs caused by sediment and other noise resulting from the camera's very large range window (Fig. 4B).

We can see the direct impact of these detection errors on our tracking and counting metrics in Fig. 4C–H. At both locations, we see that abundant FPs cause very low CLR Precision (Fig. 4C), ID Precision (Fig. 4D), and Detection Precision (Fig. 4G), negatively impacting MOTA, IDF1, and HOTA, respectively. Interestingly, while FPs cause overcounting errors at KC, the majority of counting errors at EL are actually due to undercounting (Fig. 4E). From manual inspection we diagnosed that this is often caused by TP tracks merging with hallucinated detections, causing the loss of a track before it has the chance to cross the counting line. These challenges demonstrate why counting, in addition to detection and tracking, is an important metric for researchers to consider.

Our baseline results indicate that (1) there are indeed generalization challenges in this domain, (2) they appear to be largely caused by location-specific environmental changes, and (3) these challenges affect all three tasks of interest: detection, tracking, and counting. Further, the low counting error on the validation set indicates that our overall tracking-by-detection approach is feasible at in-distribution locations and that the key challenges are in generalization. To verify this, in the next section we perform ablation studies and examine the upper-bound generalization capabilities of the proposed tracking-by-detection baseline, and use these results to motivate an improved baseline in Sec. 5.3.

5.2 Ablation Study and Generalization Upper Bounds

To evaluate generalization potential, we perform a set of "reversed" ablation studies and compare our baseline results with three different upper bounds that utilize different types of ground-truth information (Fig. 5): 1. Perfect Tracker Generalization. We use the detections from our baseline detector, but we fit the tracker hyperparameters directly to each test set based on counting performance. This results in five different trackers, one for each test set, with the best possible tracker parameter settings at each location. This gives an upper bound for the counting performance of a multiple-object tracker which can perfectly generalize to new locations.

2. Perfect Detector. This model takes *ground-truth* detections as input. The tracker hyperparameters are fit to counting performance on the validation set as in our baseline. This gives an upper bound for the counting performance of our tracker when given perfect detections.

3. Oracle. This model combines the first two upper bounds. It takes groundtruth detections as input and fits the tracker hyperparameters to each test set. This gives us the overall upper bound on counting performance for our baseline approach given perfect detections and perfect tracker generalization.

Analysis. The most apparent result is the very strong performance at all locations given a Perfect Detector. In most cases, with perfect detections our tracker achieves near-Oracle performance and generalizes well without modification. Only one location (EL) shows significant further improvements in counting error in the Oracle method compared to the Perfect Detector method. Meanwhile, Perfect Tracker Generalization does improve performance in most cases, but not as much as the Perfect Detector. This indicates that the proposed motion-based tracking approach is indeed feasible but is dependent upon a strong detector with strong generalization capabilities. Therefore the most effective improvement to overall system-generalization performance appears to be improving the generalization capabilities of the detector, which we address in the next section.

5.3 Baseline++

Given the results from our upper bound analysis, we implemented an improved baseline method, "Baseline++," with the primary goal of improving objectdetection generalization performance. We noticed that the background (1) varies significantly across locations and (2) occludes fish (see Fig. 2 for frame examples). Thus, we appended two additional channels to our image input: (1) a background-subtracted version of each frame, where the background for each clip is obtained by averaging all frames, and (2) the difference between each background-subtracted frame and its preceding frame, to capture motion information. Example frames illustrating these transformations are included in the supplemental material. We trained a new detector and tracker with this input in the same way as the baseline model. Results are shown in Figs. 3, 4, and 5. **Analysis.** In Fig. 3, we see that our Baseline++ method leads to modest improvements on our validation set (+1.6 AP50, -1.7% nMAE), but significant improvements in generalization performance (e.g. -40.2% nMAE at KC). In Fig. 4, we dissect these improvements by looking at: (A)-(B) two example frames from KC and EL, and (C)-(H) breakdowns of tracking and counting metrics across all locations. We see the efficacy of simple background subtraction as a generalization mechanism, helping significantly reduce the number of FPs and FNs

in the example KC and EL frames as well as at all test locations. We also see evidence of some outstanding issues: one FN remains in the KC example due to a small, stationary fish that appears to have been removed by the backgroundsubtraction routine, and one FP remains in the EL example due to noise near the transducer. IDR now lags behind IDP at all locations except for EL, indicating that most remaining tracking problems are track splits causing undercounting errors. These trends are also indicated in the HOTA decompositions (Fig. 4F– H), which show that DetRe now lags behind DetPr at all locations except EL, and while AssRe has improved (i.e. track splits have been reduced), it is still lower than AssPr.

6 Conclusions

We present the Caltech Fish Counting Dataset, a natural-world sonar video dataset that allows us to study object detection, multiple-object tracking, and counting under challenging real-world domain shifts. Due to the visual qualities of the source domain of river-based sonar, the dataset poses challenges to existing methods developed primarily for urban environments and provides a benchmark for video-based counting in the wild, a task that lacks supporting benchmarks.

Our experiments show that there is still significant room for improvement in the generalization performance of tracking and counting algorithms. There are also opportunities to improve tracker generalization and making trackers more robust to noisy detections. Robust algorithms that work across the range of generalization challenges in CFC will certainly be impactful in other domains, and we hope that our dataset will provide a useful testing ground for the computer vision community to push forward progress on these tasks. High-performing methods would enable sonar-based fish counting to scale globally and have real-world impact in managing some of the world's most sensitive and valuable ecosystems.

In the future, the dataset will be expanded to include additional input formats, locations, and species. One additional path forward for the community is to explore the impact of utilizing unlabeled data for unsupervised domain adaptation or self-supervised pretraining on unseen locations, which we plan to make possible by releasing additional unlabeled data from all test locations.

Acknowledgements. We are grateful to AWS for a gift to Trout Unlimited (TU) that supported data annotations, computational and storage costs, and to the Resnick Sustainability Institute at Caltech for funding to SB and PP. An NSF Fellowship supported SB. JK, SD, and EY volunteered their time. GVH was supported by the Macaulay Library at Cornell University. For collecting the dataset, and for feedback, encouragement, and moral support, we are grateful to: George Pess and Oleksandr Stefankiv (Northwest Fisheries Science Center); James Miller, Carl Pfisterer, Dawn Wilburn, Brandon Key, Suzanne Maxwell, Gregory Buck, April Faulkner, and Jordan Head (Alaska Department of Fish and Game); Dave Kajtaniak and Michael Sparkman (California Department of Fish and Wildlife); Dean Finnerty (TU's Wild Steelhead Project); and Keith Denton, Mike McHenry, and the Lower Elwha Klallam Tribe.

References

- Ahumada, J.A., Fegraus, E., Birch, T., Flores, N., Kays, R., O'Brien, T.G., Palmer, J., Schuttler, S., Zhao, J.Y., Jetz, W., et al.: Wildlife insights: A platform to maximize the potential of camera trap and other passive sensor wildlife data for the planet. Environmental Conservation 47(1), 1–6 (2020)
- Anton, V., Hartley, S., Geldenhuis, A., Wittmer, H.U.: Monitoring the mammalian fauna of urban areas using remote cameras and citizen science. Journal of Urban Ecology 4(1), juy002 (2018)
- Arac, A., Zhao, P., Dobkin, B.H., Carmichael, S.T., Golshani, P.: Deepbehavior: A deep learning toolbox for automated analysis of animal and human behavior imaging data. Frontiers in systems neuroscience 13, 20 (2019)
- 4. Arteta, C., Lempitsky, V., Zisserman, A.: Counting in the wild. In: European conference on computer vision. pp. 483–498. Springer (2016)
- Austin, C.P., Battey, J.F., Bradley, A., Bucan, M., Capecchi, M., Collins, F.S., Dove, W.F., Duyk, G., Dymecki, S., Eppig, J.T., et al.: The knockout mouse project. Nature genetics 36(9), 921 (2004)
- Australian Institute of Marine Science (AIMS) and University of Western Australia (UWA) and Curtin University: Ozfish dataset - machine learning dataset for baited remote underwater video stations (2019)
- Bai, H., Cheng, W., Chu, P., Liu, J., Zhang, K., Ling, H.: Gmot-40: A benchmark for generic multiple object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6719–6728 (2021)
- Beery, S., Agarwal, A., Cole, E., Birodkar, V.: The iwildcam 2021 competition dataset. arXiv preprint arXiv:2105.03494 (2021)
- Beery, S., Van Horn, G., Mac Aodha, O., Perona, P.: The iwildcam 2018 challenge dataset. arXiv preprint arXiv:1904.05986 (2019)
- Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: Proceedings of the European conference on computer vision (ECCV). pp. 456–473 (2018)
- Berg, T., Liu, J., Woo Lee, S., Alexander, M.L., Jacobs, D.W., Belhumeur, P.N.: Birdsnap: Large-scale fine-grained visual categorization of birds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2011– 2018 (2014)
- Berger-Wolf, T.Y., Rubenstein, D.I., Stewart, C.V., Holmberg, J.A., Parham, J., Menon, S., Crall, J., Van Oast, J., Kiciman, E., Joppa, L.: Wildbook: Crowdsourcing, computer vision, and data science for conservation. arXiv preprint arXiv:1710.08880 (2017)
- Bernardin, K., Stiefelhagen, R.: Evaluating multiple object tracking performance: the clear mot metrics. EURASIP Journal on Image and Video Processing 2008, 1–10 (2008)
- Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B.: Simple online and realtime tracking. In: 2016 IEEE international conference on image processing (ICIP). pp. 3464–3468. IEEE (2016)
- Blanchard, G., Lee, G., Scott, C.: Generalizing from several related classification tasks to a new unlabeled sample. Advances in neural information processing systems 24 (2011)
- 16. Boenisch, F., Rosemann, B., Wild, B., Dormagen, D., Wario, F., Landgraf, T.: Tracking all members of a honey bee colony over their lifetime using learned models of correspondence. Frontiers in Robotics and AI 5, 35 (2018)

- 16 J. Kay et al.
- Bogue, M.A., Philip, V.M., Walton, D.O., Grubb, S.C., Dunn, M.H., Kolishovski, G., Emerson, J., Mukherjee, G., Stearns, T., He, H., et al.: Mouse phenome database: a data repository and analysis suite for curated primary mouse phenotype data. Nucleic acids research 48(D1), D716–D723 (2020)
- Bolkensteyn, D.: dbolkensteyn/vatic.js (May 2020), https://github.com/ dbolkensteyn/vatic.js, original-date: 2016-11-23T12:39:07Z
- Bondi, E., Jain, R., Aggrawal, P., Anand, S., Hannaford, R., Kapoor, A., Piavis, J., Shah, S., Joppa, L., Dilkina, B., Tambe, M.: Birdsai: A dataset for detection and tracking in aerial thermal infrared videos. In: WACV (2020)
- Boom, B.J., He, J., Palazzo, S., Huang, P.X., Beyan, C., Chou, H.M., Lin, F.P., Spampinato, C., Fisher, R.B.: A research tool for long-term and continuous analysis of fish assemblage in coral-reefs using underwater camera footage. Ecological Informatics 23, 83–97 (2014)
- Bozek, K., Hebert, L., Mikheyev, A.S., Stephens, G.J.: Towards dense object tracking in a 2d honeybee hive. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4185–4193 (2018)
- 22. Brandt, M., Tucker, C., Kariryaa, A., Rasmussen, K., Abel, C., Small, J., Chave, J., Rasmussen, L., Hiernaux, P., Diouf, A., Kergoat, L., Mertz, O., Igel, C., Gieseke, F., Schöning, J., Li, S., Melocik, K., Meyer, J., Sinno, S., Romero, E., Glennie, E., Montagu, A., Dendoncker, M., Fensholt, R.: An unexpectedly large count of trees in the west african sahara and sahel (2020). https://doi.org/10.3334/ORNLDAAC/1832
- Bruslund Haurum, J., Karpova, A., Pedersen, M., Hein Bengtson, S., Moeslund, T.B.: Re-identification of zebrafish using metric learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops. pp. 1–11 (2020)
- Bui, N., Yi, H., Cho, J.: A vehicle counts by class framework using distinguished regions tracking at multiple intersections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 578–579 (2020)
- Chan, A.B., Liang, Z.S.J., Vasconcelos, N.: Privacy preserving crowd monitoring: Counting people without people models or tracking. In: 2008 IEEE conference on computer vision and pattern recognition. pp. 1–7. IEEE (2008)
- Change Loy, C., Gong, S., Xiang, T.: From semi-supervised to transfer counting of crowds. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2256–2263 (2013)
- Chen, K., Loy, C.C., Gong, S., Xiang, T.: Feature mining for localised crowd counting. In: Bmvc. vol. 1, p. 3 (2012)
- Chen, Y., Li, W., Gool, L.V.: Road: Reality oriented adaptation for semantic segmentation of urban scenes. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7892–7901 (2018)
- Ciaparrone, G., Sánchez, F.L., Tabik, S., Troiano, L., Tagliaferri, R., Herrera, F.: Deep learning in video multi-object tracking: A survey. Neurocomputing **381**, 61–88 (2020)
- Csurka, G.: Domain adaptation for visual applications: A comprehensive survey. arXiv preprint arXiv:1702.05374 (2017)
- Cutter, G., Stierhoff, K., Zeng, J.: Automated detection of rockfish in unconstrained underwater videos using haar cascades and a new image dataset: labeled fishes in the wild. In: 2015 IEEE Winter Applications and Computer Vision Workshops. pp. 57–62. IEEE (2015)

- 32. van Dam, E.A., van der Harst, J.E., ter Braak, C.J., Tegelenbosch, R.A., Spruijt, B.M., Noldus, L.P.: An automated system for the recognition of various specific rat behaviours. Journal of neuroscience methods **218**(2), 214–224 (2013)
- Dave, A., Khurana, T., Tokmakov, P., Schmid, C., Ramanan, D.: Tao: A largescale benchmark for tracking any object. In: European conference on computer vision. pp. 436–454. Springer (2020)
- Dendorfer, P., Rezatofighi, H., Milan, A., Shi, J., Cremers, D., Reid, I., Roth, S., Schindler, K., Leal-Taixé, L.: Mot20: A benchmark for multi object tracking in crowded scenes. arXiv preprint arXiv:2003.09003 (2020)
- Ditria, E.M., Connolly, R.M., Jinks, E.L., Lopez-Marcano, S.: Annotated video footage for automated identification and counting of fish in unconstrained seagrass habitats. Frontiers in Marine Science 8, 160 (2021)
- Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International journal of computer vision 88(2), 303–338 (2010)
- Eyjolfsdottir, E., Branson, K., Yue, Y., Perona, P.: Learning recurrent representations for hierarchical behavior modeling. arXiv preprint arXiv:1611.00094 (2016)
- Eyjolfsdottir, E., Branson, S., Burgos-Artizzu, X.P., Hoopfer, E.D., Schor, J., Anderson, D.J., Perona, P.: Detecting social actions of fruit flies. In: European Conference on Computer Vision. pp. 772–787. Springer (2014)
- 39. Fang, C., Xu, Y., Rockmore, D.N.: Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1657–1664 (2013)
- Fang, Y., Zhan, B., Cai, W., Gao, S., Hu, B.: Locality-constrained spatial transformer network for video crowd counting. In: 2019 IEEE International Conference on Multimedia and Expo (ICME). pp. 814–819. IEEE (2019)
- 41. Fennell, M., Beirne, C., Burton, A.C.: Use of object detection in camera trap image identification: assessing a method to rapidly and accurately classify human and animal detections for research and application in recreation ecology. bioRxiv (2022). https://doi.org/10.1101/2022.01.14.476404, https://www.biorxiv.org/ content/early/2022/01/21/2022.01.14.476404
- Fernandes, A.F.A., Dórea, J.R.R., Rosa, G.J.d.M.: Image analysis and computer vision applications in animal sciences: an overview. Frontiers in Veterinary Science 7, 551269 (2020)
- 43. G. Rahr: Why Protect Salmon, https://www.wildsalmoncenter.org/ why-protect-salmon/
- Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3354–3361. IEEE (2012)
- Gemert, J.C.v., Verschoor, C.R., Mettes, P., Epema, K., Koh, L.P., Wich, S.: Nature conservation drones for automatic localization and counting of animals. In: European Conference on Computer Vision. pp. 255–270. Springer (2014)
- Geuther, B.Q., Deats, S.P., Fox, K.J., Murray, S.A., Braun, R.E., White, J.K., Chesler, E.J., Lutz, C.M., Kumar, V.: Robust mouse tracking in complex environments using neural networks. Communications biology 2(1), 1–11 (2019)
- 47. Geuther, B.Q., Peer, A., He, H., Sabnis, G., Philip, V.M., Kumar, V.: Action detection using a neural network elucidates the genetics of mouse grooming behavior. Elife 10, e63207 (2021)
- Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. arXiv preprint arXiv:2007.01434 (2020)

- 18 J. Kay et al.
- 49. Hoffman, J., Wang, D., Yu, F., Darrell, T.: Fcns in the wild: Pixel-level adversarial and constraint-based adaptation (2016)
- Holmberg, J., Norman, B., Arzoumanian, Z.: Estimating population size, structure, and residency time for whale sharks rhincodon typus through collaborative photo-identification. Endangered Species Research 7(1), 39–53 (2009)
- 51. Hong, W., Kennedy, A., Burgos-Artizzu, X.P., Zelikowsky, M., Navonne, S.G., Perona, P., Anderson, D.J.: Automated measurement of mouse social behaviors using depth sensing, video tracking, and machine learning. Proceedings of the National Academy of Sciences **112**(38), E5351–E5360 (2015)
- 52. Hornakova, A., Henschel, R., Rosenhahn, B., Swoboda, P.: Lifted disjoint paths with application in multiple object tracking. In: International Conference on Machine Learning. pp. 4364–4375. PMLR (2020)
- 53. Hsieh, M.R., Lin, Y.L., Hsu, W.H.: Drone-based object counting by spatially regularized regional proposal network. In: Proceedings of the IEEE international conference on computer vision. pp. 4145–4153 (2017)
- 54. Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S., et al.: Speed/accuracy trade-offs for modern convolutional object detectors. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7310–7311 (2017)
- 55. Idrees, H., Tayyab, M., Athrey, K., Zhang, D., Al-Maadeed, S., Rajpoot, N., Shah, M.: Composition loss for counting, density map estimation and localization in dense crowds. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 532–546 (2018)
- 56. Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., NanoCode012, Kwon, Y., TaoXie, Fang, J., imyhxy, Michael, K., Lorna, V, A., Montes, D., Nadar, J., Laughing, tkianai, yxNONG, Skalski, P., Wang, Z., Hogan, A., Fati, C., Mammana, L., AlexWang1900, Patel, D., Yiwei, D., You, F., Hajek, J., Diaconu, L., Minh, M.T.: ultralytics/yolov5: v6.1 - TensorRT, TensorFlow Edge TPU and OpenVINO Export and Inference (Feb 2022). https://doi.org/10.5281/zenodo.6222936, https://doi.org/10. 5281/zenodo.6222936
- 57. Jones, F.M., Allen, C., Arteta, C., Arthur, J., Black, C., Emmerson, L.M., Freeman, R., Hines, G., Lintott, C.J., Macháčková, Z., et al.: Time-lapse imagery and volunteer classifications from the zooniverse penguin watch project. Scientific data 5(1), 1–13 (2018)
- 58. Kalman, R.E.:А New Approach to Linear Filtering and Prediction Problems. Journal Basic Engineering **82**(1), 35 of 451960). https://doi.org/10.1115/1.3662552, (Mar https:// asmedigitalcollection.asme.org/fluidsengineering/article/82/1/35/ 397706/A-New-Approach-to-Linear-Filtering-and-Prediction
- Kamenetsky, D., Sherrah, J.: Aerial car detection and urban understanding. In: 2015 International Conference on Digital Image Computing: Techniques and Applications (DICTA). pp. 1–8. IEEE (2015)
- Kellenberger, B., Marcos, D., Tuia, D.: Detecting mammals in uav images: Best practices to address a substantially imbalanced dataset with deep learning. Remote sensing of environment 216, 139–153 (2018)
- Key, B., Miller, J., Huang, J.: Operational plan: Kenai river chinook salmon sonar assessment at river mile 13.7, 2020–2022 (2020)
- Kocamaz, M.K., Gong, J., Pires, B.R.: Vision-based counting of pedestrians and cyclists. In: 2016 IEEE winter conference on applications of computer vision (WACV). pp. 1–8. IEEE (2016)

- Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., et al.: Wilds: A benchmark of inthe-wild distribution shifts. In: International Conference on Machine Learning. pp. 5637–5664. PMLR (2021)
- 64. Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., et al.: Wilds: A benchmark of inthe-wild distribution shifts. In: International Conference on Machine Learning. pp. 5637–5664. PMLR (2021)
- Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: 2011 International conference on computer vision. pp. 2556–2563. IEEE (2011)
- 66. Kulits, P., Wall, J., Bedetti, A., Henley, M., Beery, S.: Elephantbook: A semiautomated human-in-the-loop system for elephant re-identification. In: ACM SIG-CAS Conference on Computing and Sustainable Societies. pp. 88–98 (2021)
- Kumar, N., Belhumeur, P.N., Biswas, A., Jacobs, D.W., Kress, W.J., Lopez, I.C., Soares, J.V.: Leafsnap: A computer vision system for automatic plant species identification. In: European conference on computer vision. pp. 502–516. Springer (2012)
- Leal-Taixé, L., Milan, A., Reid, I., Roth, S., Schindler, K.: Motchallenge 2015: Towards a benchmark for multi-target tracking. arXiv preprint arXiv:1504.01942 (2015)
- 69. Li, S., Li, J., Lin, W., Tang, H.: Amur tiger re-identification in the wild. arXiv e-prints pp. arXiv-1906 (2019)
- Liu, L., Lu, H., Cao, Z., Xiao, Y.: Counting Fish in Sonar Images. In: 2018 25th IEEE International Conference on Image Processing (ICIP). pp. 3189–3193 (Oct 2018). https://doi.org/10.1109/ICIP.2018.8451154, iSSN: 2381-8549
- Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixé, L., Leibe, B.: Hota: A higher order metric for evaluating multi-object tracking. International journal of computer vision 129(2), 548–578 (2021)
- Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., Kim, T.K.: Multiple object tracking: A literature review. Artificial Intelligence 293, 103448 (2021)
- Ma, Z., Chan, A.B.: Crossing the line: Crowd counting by integer programming with local features. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2539–2546 (2013)
- Mandal, V., Adu-Gyamfi, Y.: Object detection and tracking algorithms for vehicle counting: a comparative analysis. Journal of Big Data Analytics in Transportation 2(3), 251–261 (2020)
- 75. Marstaller, J., Tausch, F., Stock, S.: Deepbees-building and scaling convolutional neuronal nets for fast and large-scale visual monitoring of bee hives. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 0–0 (2019)
- 76. Mathis, M.W., Mathis, A.: Deep learning tools for the measurement of animal behavior in neuroscience. Current opinion in neurobiology **60**, 1–11 (2020)
- McCann, E., Li, L., Pangle, K., Johnson, N., Eickholt, J.: An underwater observation dataset for fish classification and fishery assessment. Scientific data 5(1), 1–8 (2018)
- Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: Mot16: A benchmark for multi-object tracking. arXiv preprint arXiv:1603.00831 (2016)
- Moranduzzo, T., Melgani, F.: Automatic car counting method for unmanned aerial vehicle images. IEEE Transactions on Geoscience and Remote Sensing 52(3), 1635–1647 (2013)

- 20 J. Kay et al.
- Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: International Conference on Machine Learning. pp. 10–18. PMLR (2013)
- Naphade, M., Wang, S., Anastasiu, D.C., Tang, Z., Chang, M.C., Yang, X., Yao, Y., Zheng, L., Chakraborty, P., Lopez, C.E., et al.: The 5th ai city challenge. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4263–4273 (2021)
- Nilsback, M.E., Zisserman, A.: A visual vocabulary for flower classification. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 1447–1454. IEEE (2006)
- Norouzzadeh, M.S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M.S., Packer, C., Clune, J.: Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. Proceedings of the National Academy of Sciences 115(25), E5716–E5725 (2018)
- Onoro-Rubio, D., López-Sastre, R.J.: Towards perspective-free object counting with deep learning. In: European conference on computer vision. pp. 615–629. Springer (2016)
- Pardo, L.E., Bombaci, S., Huebner, S.E., Somers, M.J., Fritz, H., Downs, C., Guthmann, A., Hetem, R.S., Keith, M., Roux, A.I., et al.: Snapshot safari: A large-scale collaborative to monitor africa's remarkable biodiversity. South African Journal of Science 117(1-2), 1–4 (2021)
- Parham, J.R., Crall, J., Stewart, C., Berger-Wolf, T., Rubenstein, D.: Animal population censusing at scale with citizen science and photographic identification. In: 2017 AAAI Spring Symposium Series (2017)
- Pedersen, M., Haurum, J.B., Bengtson, S.H., Moeslund, T.B.: 3d-zef: A 3d zebrafish tracking benchmark dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2426–2436 (2020)
- Revaud, J., Weinzaepfel, P., Harchaoui, Z., Schmid, C.: Deepmatching: Hierarchical deformable dense matching. International Journal of Computer Vision 120(3), 300–323 (2016)
- Rey, N., Volpi, M., Joost, S., Tuia, D.: Detecting animals in african savanna with uavs and the crowds. Remote Sensing of Environment 200, 341–351 (2017)
- Richards, B.L., Drazen, J.C., Virginia Moriwake, V.: Hawai'i deep-7 bottomfish training and validation image dataset: Noaa pacific islands fisheries science center botcam stereo-video (2014)
- Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C.: Performance measures and a data set for multi-target, multi-camera tracking. In: European conference on computer vision. pp. 17–35. Springer (2016)
- Rodriguez, I.F., Megret, R., Acuna, E., Agosto-Rivera, J.L., Giray, T.: Recognition of pollen-bearing bees from video using convolutional neural network. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 314–322. IEEE (2018)
- Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: European conference on computer vision. pp. 213–226. Springer (2010)
- 94. Saleh, A., Laradji, I.H., Konovalov, D.A., Bradley, M., Vazquez, D., Sheaves, M.: A realistic fish-habitat dataset to evaluate algorithms for underwater visual analysis. Scientific Reports 10(1), 1–10 (2020)
- 95. Sam, D.B., Peri, S.V., Sundararaman, M.N., Kamath, A., Radhakrishnan, V.B.: Locate, size and count: Accurately resolving people in dense crowds via detection. IEEE transactions on pattern analysis and machine intelligence (2020)

21

- Schneider, S., Zhuang, A.: Counting fish and dolphins in sonar images using deep learning. arXiv preprint arXiv:2007.12808 (2020)
- 97. Shao, W., Kawakami, R., Yoshihashi, R., You, S., Kawase, H., Naemura, T.: Cattle detection and counting in uav images based on convolutional neural networks. International Journal of Remote Sensing 41(1), 31–52 (2020)
- Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012)
- 99. Stierhoff, K., Cutter, G.: Rockfish (sebastes spp.) training and validation image dataset: Noaa southwest fisheries science center remotely operated vehicle (rov) digital still images (2013)
- 100. Sun, J.J., Karigo, T., Chakraborty, D., Mohanty, S.P., Wild, B., Sun, Q., Chen, C., Anderson, D.J., Perona, P., Yue, Y., et al.: The multi-agent behavior dataset: Mouse dyadic social interactions. arXiv preprint arXiv:2104.02710 (2021)
- 101. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2446–2454 (2020)
- 102. Swanson, A., Kosmala, M., Lintott, C., Simpson, R., Smith, A., Packer, C.: Snapshot serengeti, high-frequency annotated camera trap images of 40 mammalian species in an african savanna. Scientific data 2(1), 1–14 (2015)
- 103. Tabak, M.A., Norouzzadeh, M.S., Wolfson, D.W., Sweeney, S.J., VerCauteren, K.C., Snow, N.P., Halseth, J.M., Di Salvo, P.A., Lewis, J.S., White, M.D., et al.: Machine learning to classify animal species in camera trap images: Applications in ecology. Methods in Ecology and Evolution 10(4), 585–590 (2019)
- 104. The Nature Conservancy: Channel islands camera traps 1.0 (2021)
- 105. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: CVPR 2011. pp. 1521–1528. IEEE (2011)
- 106. Tuia, D., Kellenberger, B., Beery, S., Costelloe, B.R., Zuffi, S., Risse, B., Mathis, A., Mathis, M.W., van Langevelde, F., Burghardt, T., et al.: Perspectives in machine learning for wildlife conservation. Nature Communications 13(1), 1–15 (2022)
- 107. Van Horn, G., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeirotis, P., Perona, P., Belongie, S.: Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 595– 604 (2015)
- 108. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8769–8778 (2018)
- 109. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
- 110. Wang, Q., Gao, J., Lin, W., Li, X.: Nwpu-crowd: A large-scale benchmark for crowd counting and localization. IEEE transactions on pattern analysis and machine intelligence 43(6), 2141–2149 (2020)
- Weinland, D., Ronfard, R., Boyer, E.: Free viewpoint action recognition using motion history volumes. Computer vision and image understanding 104(2-3), 249–257 (2006)
- 112. Weinstein, B.G., Marconi, S., Bohlman, S.A., Zare, A., Singh, A., Graves, S.J., White, E.P.: A remote sensing derived data set of 100 million individual tree crowns for the national ecological observatory network. Elife **10**, e62922 (2021)

- J. Kay et al.
- 113. Wen, L., Du, D., Cai, Z., Lei, Z., Chang, M.C., Qi, H., Lim, J., Yang, M.H., Lyu, S.: Ua-detrac: A new benchmark and protocol for multi-object detection and tracking. Computer Vision and Image Understanding 193, 102907 (2020)
- 114. Wen, L., Du, D., Zhu, P., Hu, Q., Wang, Q., Bo, L., Lyu, S.: Detection, tracking, and counting meets drones in crowds: A benchmark. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7812– 7821 (2021)
- 115. Wen, L., Du, D., Zhu, P., Hu, Q., Wang, Q., Bo, L., Lyu, S.: Detection, tracking, and counting meets drones in crowds: A benchmark. In: CVPR (2021)
- Wojke, N., Bewley, A., Paulus, D.: Simple online and realtime tracking with a deep association metric. In: 2017 IEEE international conference on image processing (ICIP). pp. 3645–3649. IEEE (2017)
- 117. Wu, Z., Fuller, N., Theriault, D., Betke, M.: A thermal infrared video benchmark for visual analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 201–208 (2014)
- 118. Ye, N., Li, K., Bai, H., Yu, R., Hong, L., Zhou, F., Li, Z., Zhu, J.: Ood-bench: Quantifying and understanding two dimensions of out-of-distribution generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7947–7958 (2022)
- 119. Yousif, H., Kays, R., He, Z.: Dynamic programming selection of object proposals for sequence-level animal species classification in the wild. IEEE Transactions on Circuits and Systems for Video Technology (2019)
- 120. Yu, F., Xian, W., Chen, Y., Liu, F., Liao, M., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving video database with scalable annotation tooling. arXiv preprint arXiv:1805.04687 2(5), 6 (2018)
- 121. Zhang, C., Kang, K., Li, H., Wang, X., Xie, R., Yang, X.: Data-driven crowd understanding: A baseline for a large-scale crowd dataset. IEEE Transactions on Multimedia 18(6), 1048–1061 (2016)
- 122. Zhang, C., Li, H., Wang, X., Yang, X.: Cross-scene crowd counting via deep convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 833–841 (2015)
- 123. Zhang, S., Wu, G., Costeira, J.P., Moura, J.M.: Fcn-rlstm: Deep spatio-temporal neural networks for vehicle counting in city cameras. In: Proceedings of the IEEE international conference on computer vision. pp. 3667–3676 (2017)
- 124. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 2039–2049 (2017)
- 125. Zhang, Y., Sun, P., Jiang, Y., Yu, D., Yuan, Z., Luo, P., Liu, W., Wang, X.: Bytetrack: Multi-object tracking by associating every detection box. arXiv preprint arXiv:2110.06864 (2021)
- 126. Zhang, Y., Wang, C., Wang, X., Zeng, W., Liu, W.: Fairmot: On the fairness of detection and re-identification in multiple object tracking. arXiv preprint arXiv:2004.01888 (2020)
- 127. Zhang, Y., Zhou, D., Chen, S., Gao, S., Ma, Y.: Single-image crowd counting via multi-column convolutional neural network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 589–597 (2016)
- 128. Zhang, Z., He, Z., Cao, G., Cao, W.: Animal detection from highly cluttered natural scenes using spatiotemporal object region proposals and patch verification. IEEE Transactions on Multimedia 18(10), 2079–2092 (2016)

- Zhao, Z., Li, H., Zhao, R., Wang, X.: Crossing-line crowd counting with two-phase deep neural networks. In: European Conference on Computer Vision. pp. 712–726. Springer (2016)
- Zheng, Z., Yang, X., Yu, Z., Zheng, L., Yang, Y., Kautz, J.: Joint discriminative and generative learning for person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2138– 2147 (2019)
- 131. Zhou, K., Liu, Z., Qiao, Y., Xiang, T., Loy, C.C.: Domain generalization in vision: A survey (2021)
- 132. Zhou, Y., Yu, H., Wu, J., Cui, Z., Zhang, F.: Fish behavior analysis based on computer vision: A survey. In: International Conference of Pioneering Computer Scientists, Engineers and Educators. pp. 130–141. Springer (2019)
- 133. Zhu, P., Peng, T., Du, D., Yu, H., Zhang, L., Hu, Q.: Graph regularized flow attention network for video animal counting from drones. IEEE Transactions on Image Processing (2021)
- 134. Zou, Z., Shi, Z., Guo, Y., Ye, J.: Object detection in 20 years: A survey. arXiv preprint arXiv:1905.05055 (2019)