
IMITATION-BASED ACTIVE CAMERA CONTROL WITH DEEP CONVOLUTIONAL NEURAL NETWORK

THIS PAPER IS A PREPRINT OF A PAPER SUBMITTED TO AND ACCEPTED FOR PUBLICATION IN PROCEEDINGS OF IEEE INTERNATIONAL CONFERENCE ON IMAGE PROCESSING, APPLICATIONS, AND SYSTEMS - IPAS2020

Christos Kyrkou*
KIOS Research and Innovation Center of Excellence
University of Cyprus
1 Panepistimiou Avenue, Nicosia Cyprus
{kyrkou.christos}@ucy.ac.cy

November 27, 2020

ABSTRACT

The increasing need for automated visual monitoring and control for applications such as smart camera surveillance, traffic monitoring, and intelligent environments, necessitates the improvement of methods for visual active monitoring. Traditionally, the active monitoring task has been handled through a pipeline of modules such as detection, filtering, and control. In this paper we frame active visual monitoring as an imitation learning problem to be solved in a supervised manner using deep learning, to go directly from visual information to camera movement in order to provide a satisfactory solution by combining computer vision and control. A deep convolutional neural network is trained end-to-end as the camera controller that learns the entire processing pipeline needed to control a camera to follow multiple targets and also estimate their density from a single image. Experimental results indicate that the proposed solution is robust to varying conditions and is able to achieve better monitoring performance both in terms of number of targets monitored as well as in monitoring time than traditional approaches, while reaching up to 25 FPS. Thus making it a practical and affordable solution for multi-target active monitoring in surveillance and smart-environment applications.

Keywords Real-Time Active Vision, Smart Camera, Deep Learning, End-to-End

1 Introduction

Active vision systems (i.e., movable cameras with controllable parameters such as pan and tilt) have received much attention in recent years due to their extended coverage, flexibility, cost-efficiency compared to static vision systems [1]. Active cameras can be used to track targets (i.e. follow them) and can provide continuous monitoring of an area reliably and robustly and are increasingly being used for various applications ranging from surveillance[2] to intelligent interactive environments[3]. However, there is a limited number of cameras that human operator can monitor and control hence, there is a need for for automated, robust, and reliable systems for active vision. In addition, there are additional challenges for systems operating with battery limitations or in remote locations and with requirements for rapid deployment in temporary installations necessitating reduced computational cost and simpler active vision systems.

Existing approaches for automated active vision decompose the active control problem into separate modules, such as detection, tracking, and control [4], and employ different algorithms at each step. Such examples include motion detection, background modelling/subtraction, and lastly tracking by detection. The former two are widely used in static camera scenarios [5] due to their relative computational efficiency however, for active cameras that move and exhibit constant background change the latter methods are preferred since they are more widely applicable. Such techniques are often augmented with post-processing filters that increase the computational complexity and may also susceptible to

*ckyrkou@gmail.com, www.christoskyrkou.com

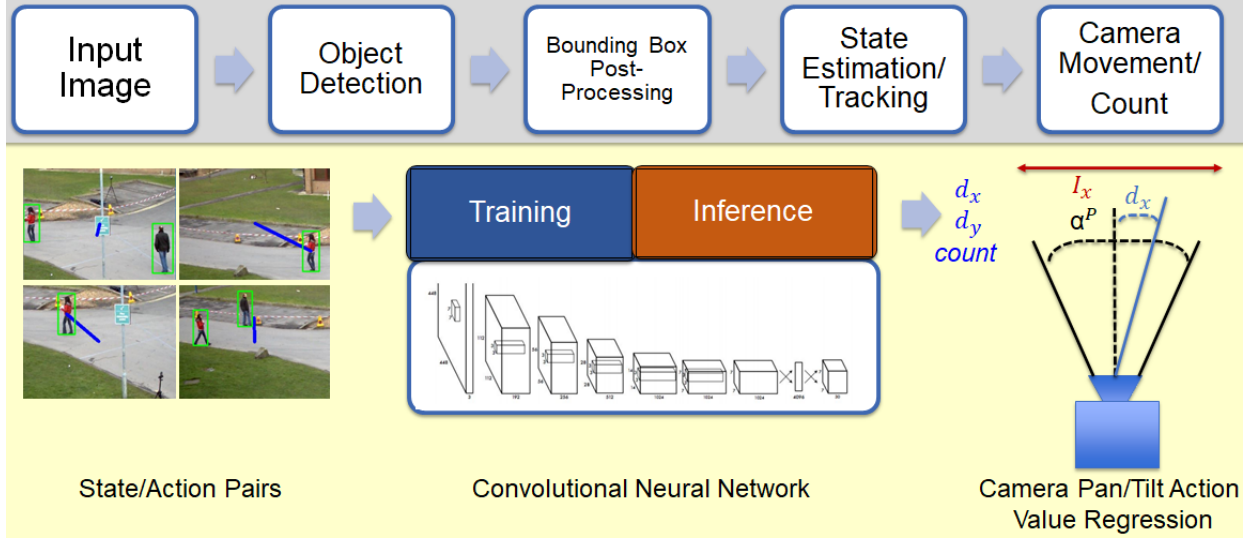


Figure 1: (top) Traditional multi-stage active monitoring vs. (bottom) End-to-end control with deep neural network (DNN). The DNN learns through examples how to control the camera to simultaneously recognize targets and keep them in its Field-of-View (FoV).

false alarms and outliers. Hence, existing methods that rely on hand-crafted features, motion models, and modeling of the camera views are not optimized for active vision scenarios.

To deal with the aforementioned challenges in this paper we propose leveraging Convolutional Neural Networks (CNN) and end-to-end imitation-learning to develop an active vision system for monitoring people in surveillance applications. To the best of our knowledge there has not yet been any attempt to deal with active tracking of multiple targets in an end-to-end way. In this work, we investigate it and present its appealing potential. In particular an Active Camera Deep Controller Network, referred to as ACDCNet, is trained in an end-to-end manner through imitation learning to associate single image features with control actions. End-to-end learning approaches allow encapsulating all the intelligence into the machine learning algorithm thus can optimizing all processing steps simultaneously and learning the features to associate with camera control for visual active tracking. In addition, it can result into smaller less complex smart camera systems.

The proposed approach has been verified by extensive experiments using simulation to replicate the motion of the camera and targets. Results indicate that the proposed deep active monitoring approach outperforms some representative methods, composed of several submodules including object detection and tracking, for following multiple targets.

2 Background and Related Work

The aim of static object tracking is to localize an object in successive video frames given an annotation the initial frame. It has gained more attention due to its relatively simpler problem setting [6]. However, the video frames are considered still and hence the approach is not applicable to active cameras where in addition to the challenges associated with moving objects and changing scene perspectives, there is also the issue of controlling the camera and optimize its position for the given vision task. Since the camera itself is constantly changing its orientation, object initialization and continuous tracking pose great challenges. Standard techniques such as foreground segmentation, motion detection or optical flow may face difficulties as the background is constantly changing.

The goal of active monitoring in contrast to a static monitoring setting is to change the camera control parameters in order to maximize a visual-task-related performance objective such as following one or more targets that are located within the field-of-view (FoV), in order to improve the overall surveillance and monitoring capabilities [7]. Active cameras have gained considerable interest in recent years[8], however, there still does not exist a generic solution. Conventional solutions for active visual monitoring tackle the problem by decomposing it into two or more sub-tasks [9], i.e., object detection typically using a machine-learning-based classifier/detector, a tracking algorithms such as Kalman filter, and a control output for the camera movement. In such case, each task is optimized individually resulting in highly complex systems with many tuning parameters. Furthermore, this leads to difficulty in obtaining real-time

performance on resource-constraint embedded camera systems. Different works have investigated the use of active cameras with one or more degrees of freedom to address the problem of monitoring an area and are summarized next.

Initial approaches such as [10], followed a master-slave approach used to track targets at a high resolution. One camera, the master, has a wide FoV and performs blob detection and uses a Kalman filter to track a target in an area and by projecting from image plane to world coordinates it controls (pan, tilt, and zoom) the other active camera, the slave, to follow a target. In contrast our goal in this work is to improve the tracking performance of a single camera agent so that it can autonomously follow the majority of targets in its FoV.

In [11] an Extended Kalman filter is used to jointly track the object position in the real world as well as estimating the intrinsic camera parameters. The filter outputs are used as inputs to two PID controllers (one for pan and one for tilt motion axis) which continuously track a moving target at a certain resolution. The focus of this work is on tracking only a single target however, and the objects are assumed to have a predetermined size.

The active camera system proposed in [3] is composed of multiple components in order to track a subject. The face of the target is identified through a face detection system and then a tracker is employed to estimate the targets motion across frames based on previous observations. An online learning approach is also used to learn the appearance model over time and reduce false detections. In addition, Gaussian Mixture Models are used to model the body movement in case the face detection fails. The final part is the controller of the pan-tilt camera which makes its decision of how to move based on where in the image the target is positioned and using a set of rules. Such multi-component system can be difficult to tune and transfer to more resource-constraint systems that monitor multiple targets.

Finally, in [12] an end-to-end active tracker is proposed for following a character in the VizDoom video game environment. It uses reinforcement learning to train an CNN with LSTM to output movement actions for following a single target. Even though they are somewhat realistic, these scenarios do not correspond to real-world use cases. Furthermore, the step-based output is not suitable for the dynamic range needed to control a pan-tilt camera.

In summary, it is evident from the literature that related works make excessive use of multiple modules composed of hand-crafted models and rules that must be tuned separately and in most cases track only a single target. While there has been considerable progression in utilizing deep learning for static camera tracking there has been relatively few works dealing with deep learning for active smart camera systems. Hence, in this work we attempt to bridge the gap between the use of active cameras with deep learning algorithms by proposing an end-to-end learning approach to simultaneously build a detector and controller for cameras with pan and tilt motion capabilities.

3 Active Camera Control with Deep Learning

3.1 Imitation Learning Approach

To solve the problem of active visual control for target following we frame the problem in an imitation learning setting to be solved in a supervised way. Under this setting the agent (learner) needs to come up with a policy whose resulting state-action behaviour matches the expert behaviour.² For this specific problem the behaviour is how to control the camera given an input frame to keep most of the targets in the field-of-view (FoV). It is assumed that expert state-action pairs are available, $(s_1, a_1), (s_2, a_2), \dots, (s_N, a_N)$ where s_i are the image frames and a_i are the associated camera actions, such that can form a training set. The expert behaviors are assumed to be i.i.d and then train a regressor network end-to-end in a supervised manner, to learn a control function f that maps observations to actions such that $a_i \simeq f(s_i)$. In doing so, the desired actions are directly regressed during inference, without having to explicitly train the network to first detect the targets (Fig. 1).

A pinhole model [8] is assumed for the camera, and the desired actions are to control its servo motors along its pan and tilt axes (denoted by a^P and a^T). Each image s_i has resolution $I_x \times I_y$. Under these circumstances the pixel distances between the image center to the center of mass of the targets is analogous to the angle that the servo motors have to move in order to position the target(s) at the center of FoV (as shown in Fig. 1) [4, 2]. In addition, the motion of the camera when tracking targets in its FoV is bounded by the half of its horizontal and vertical angles of view (denoted as θ_x, θ_y). Accordingly, the motion of the camera in the pan and tilt axes can be calculated with respect to the viewing angles. This can be done by first calculating the distance of the camera center from the target(s) center of mass for the horizontal and vertical direction (d_x, d_y) , and then use Eq. 1 and 2 to associate the pixel distance to the camera angles as also illustrated in (Fig. 1). Hence, during the learning process the objective is to map the input image to normalized pixel displacement values $(d_x/I_x, d_y/I_y)$, which will then be used to calculate the corresponding servo motor angle

²Such expert behaviour can be gathered by appropriate datasets with bounding box annotations for targets from which the camera movement can be extracted.



Figure 2: Example of augmented images from the dataset.

displacements in the pan and tilt axis. Predicting offsets instead of angles simplifies the problem and makes it easier for the network to learn. Furthermore, it decouples the learning process from camera specific parameters.

$$a^P \rightarrow \frac{d_x}{I_x} \times \frac{\theta_x}{2} \quad (1)$$

$$a^T \rightarrow \frac{d_y}{I_y} \times \frac{\theta_y}{2} \quad (2)$$

3.2 Data Collection and Training

Proper training and testing data are necessary to train a deep CNN regressor for the visual active monitoring task. To the best of the author’s knowledge there is no publicly available dataset for active vision surveillance applications with ground truth camera controls. For this reason existing wide surveillance image databases, such as the PETS2009[13], used for static tracking, are re-purposed to develop a simulation framework that allows for acquiring state-action pairs corresponding to an image frame with the action in image space needed to position the camera to the center-of-mass of the targets in the frame. These ground truth information that can be used for both training and evaluation. In addition, it allows for simulating active cameras and evaluating performance with real-world images, and capturing and storing multiple frame sequences with expert ground truth data for bounding box, density, and camera control. The simulation framework also simulates the movement of a virtual active camera with pan-tilt control with a fixed FoV that is able to move within a larger image frame.

An initial dataset of over 3100 images was generated using the simulation software and sequence from the PETS2009[13] database in order to train ACDCNet. Starting at various positions in random frames in the sequence we calculate the changes in displacement of pan and tilt axes according to Eq. 1 and 2 so that the targets will be positioned at the camera image center in the next frame. Hence, the training set effectively contains single images sampled from the sequence, paired with the corresponding pan and tilt displacement that must be applied. The pan and tilt displacement values are normalized between -1 and 1 with regards to the maximum motion of the camera in both axes. As such, it is also possible to apply any mapping function to scale the output depending on the required configuration. It is also possible to extract different resolution images with varying number of targets in each one, For this work images of 320×240 are targeted. It is worth noting that any input image size can be processed after being re-sized appropriately and from cameras with different angles of view. Examples of some of the images in the dataset are shown in Fig. 2

An important part of data preparation is balancing the data. During the data collection process depending on the motion of successive frames, the target displacement values may not change significantly in both axes. Also successive frames may have high correlation and hence little to no motion in both axes. As a result, as shown in Fig. 3 the distribution of camera motion values might be imbalanced. To ensure that the learning algorithm does not overfit to the majority of values around zero, the data in each batch is sampled to contain high as well as low values so that the network better learns when and how to move.

The whole data is split into 75% for training the CNN regression model, and 25% of the data was used for testing the final accuracy of the model as well as to compare with other methods. Furthermore augmentations are probabilistically applied to the images during training, to increase the variability and combat over-fitting. The augmentation strategy




Figure 3: Distribution of the pan and tilt values from the images in the dataset.

included some transformations on the image pixels such as blurring and sharpening, color-shifting, illumination changes as well as changes in the image orientation such as translations and horizontal flip with appropriate adjustment of the target pan and tilt values if a geometric augmentation is applied. The combinations of all these augmentations resulted in a variety of novel images that were used for training.

3.3 Active Camera Deep Controller Network (ACDCNet)

ACDCNet is trained to perform the regression tasks of estimating the camera motion displacement and the number of targets in the frame. Figure 4 shows the network architecture which is conceptually comprised of two main parts, a feature extractor made up of convolutional layers and a motion controller that summarizes the feature maps in order to produce the final output values. Although it is difficult to distinguish the parts of the network that function primarily as feature extractor and which serve as controller since the system is trained end-to-end. The input to the CNN is a 320×240 RGB image which is normalized by dividing the pixels by 255.




Figure 4: ACDCNet Architecture: Conceptually comprised of a feature extractor and a controller that condenses the high level semantic information to infer an action. Then it regresses the motion for the camera head along the pan and tilt axes as well as the number of targets in the frame.

Convolutional Feature Extractor: The layers were designed to perform feature extraction and were chosen empirically through a series of experiments that varied layer configurations. There are 7 major blocks each comprised of a convolutional layer with *leaky relu* activation with $\alpha = 0.3$ and batch-normalization layer. Furthermore, dropout is applied at the middle of the sub-network to combat overfitting with a rate of 0.2. Overall, the feature extractor is designed to be inherently computationally efficient to support use in embedded smart cameras for local control and decision making. For this reason the first 3 layers downsample the image to reduce the computational cost and have a small number of filters, and overall the filters do not exceed 128. After the first 3 layers the resolution is not reduced further to improve accuracy. The primary goal is to build a small neural network for on-board processing but alternatively transfer learning techniques using other backbone networks can also be applied.

Controller Subnetwork: The controller subnetwork is comprised of both convolutional as well as fully-connected layers. The idea behind this is that the convolutional layers will condense the information from the feature extractor; then the final convolution will convert the feature map into a vector. It is then followed by 2 fully connected layers with 100 and 50 neurons respectively to map the vector to the motion controls. There is a dropout in between the dense layers and all layers have an *elu* activation function for faster convergence. The output of the controller subnetwork is further processed through a clipped linear activation that bounds the output between $[-1, \dots, 1]$ to estimate the motion in the horizontal and vertical direction more effectively. The third output neuron that regresses the number of targets uses a ReLU function to discard negative numbers. Overall, ACDCNet has a total of $\sim 386,000$ parameters. This results in a small network which requires $\sim 4\text{MB}$, resulting in a lightweight network that can run even on low-end CPUs.

3.4 Network Training

The objective of the learning process is to regress two motion values in the pan and tilt axes, and the number of targets. Accordingly, the loss function in Eq. 3 is employed for learning the camera controls and target count. The *Keras* deep learning framework [14] with *Tensorflow* [15] running as the backend is used for the training of the CNN regressor. The network was trained using a GeForce Titan Xp, on a PC with an Intel *i7 - 7700K* processor, and 32GB of RAM. The Adam optimization method was used for training with a learning rate step-decay approach starting from an initial learning rate of 0.001, and decreasing it by a factor of 0.95 every 5 epochs. The CNN regressor is trained for 500 epochs with a batch size of 32 resulting in 3100 generated training images per epoch.

$$L = \frac{1}{N_B} \sum_{j=1}^{N_B} \left[(y_c^{true}(j) - y_c^{pred}(j))^2 + |y_{d_x}^{true}(j) - y_{d_x}^{pred}(j)| + |y_{d_y}^{true}(j) - y_{d_y}^{pred}(j)| \right], \quad (3)$$




Figure 5: Errors for the different regressed quantities for each approach.

4 Evaluation and Experimental Results

To compare the proposed method a traditional active tracking pipeline is implemented following other relevant approaches for active camera control [9]. In this paradigm an object detector is used to localize the targets and a tracker is applied to filter the detections over time. Finally based on the filter target positions the camera is moved to position the targets as close to the center as possible. The control is performed using the simulation framework as discussed in Section 3.2 and is the same for all methods.

We use the YOLO [16] framework to comparison with single-shot detectors which are considered more suitable for real-time applications and can be leveraged for real-time embedded systems. We show herein comparisons with the original smaller YOLO variant referred to as *tinyYOLO*. Comparisons are also made against the widely used SVM-HOG pedestrian detection algorithm [17] implemented in OpenCV [18] and still favored for embedded applications and especially suitable for low-power systems [2]. On top of the detectors we apply tracking with Kalman filter (denoted with $_T$ in the Figures) which is a common approach used in active tracking [8, 11]. The incorporated filtering handles bounding box associations, maintains trajectories, and handles the creation and termination of tracks.

The proposed model is evaluated experimentally in a number of ways to determine its overall performance. For the first experiment we use a test set of 500 images not used in the training process, with ground truth camera motion information in order to calculate how accurate is each method in estimating the camera motion controls for positioning targets at the center of its FoV as well as regressing the number of targets. As shown in Fig. 5 the proposed CNN achieves the lowest error across the different methods. In practice this means that it will follow the targets better since it has not have to deal with bounding box localization errors that can steer off the camera control. In particular, for the case of estimating the number of targets the other methods can be as off as 1-2 targets. Furthermore, the lower errors also mean that in cases where there is no target in its FoV the network has recognized that it should not move, and in case of many targets it attempts to follow their center of mass thus keeping the majority of them in view.

The next experiment involved evaluating all the aforementioned methods in the simulation environment and measure important metrics such as the average number of targets followed. This provides an indication of how well each approach manages to keep up with the target motion as well as how many targets it can keep in its field of view. To conduct these sequences where generated from the PETS2009[13] dataset for which the ground truth bounding box annotations is available but have not been used in the training and validation phases. These sequences provide an additional challenge as they feature different visual conditions with high density crowd densities and motion patterns. The images in the sequences are of 768×576 resolution and the virtual camera FoV is set to 320×240 for all methods so that there is margin for the camera to move and follow the targets.

The developed simulation framework outlined in section 3.2 models the camera motion based on the visual input and evaluates its performance. The output of each vision pipeline is a motion vector that will be passed to the simulator to perform the action. The camera FoV for all methods is set at the same initial position. In all cases the objective of the camera is to move in such a way as to keep the most number of targets in its FoV. We also report the ground truth motion which represents the motion of the camera when the target positions are perfectly known.

The results are presented in Fig. 6 and Fig. 7. The proposed system manages to surpass the other approaches and be close to the optimum based on the number of targets tracked over the length of the sequences. The proposed approach also manages to have a performance of ~ 25 frames per second on a CPU. Since this is a simulation the effect of processing time was not evaluated directly.




Figure 6: Average number of targets monitored over time for each approach.




Figure 7: Frames-per-Second (FPS) achieved by each approach.

However, since the inference time of the other methods is higher it is expected that in real conditions slower methods will miss even more targets. Hence, the proposed approach provides a much more efficient way of building reactive active vision systems that can also be more responsive to the target motion. Overall the performance gains between the proposed and other methods can be attributed end-to-end nature of our approach that associates lower-level-features with control actions and does not rely on bounding boxes. As a result ACDCNet manages to handle more dense targets better. Fig. 8 demonstrates some cases that cause errors in the traditional detection approaches; while Fig. 9 also demonstrates the fact that the approach can generalize in out of domain images gathered from real-world settings.

5 Conclusion

This work addresses the problem of controlling a smart camera for active vision applications (surveillance, drones, vehicles, smart environments). One of the major contributions is that the problem is tackled via end-to-end learning using deep convolutional neural networks. A CNN architecture referred to as ACDCNet is proposed that maps input images to pan and tilt motion command with the goal of keeping targets in its field-of-view. Even with single frame information the proposed network outperforms other multi-stage approaches in terms of monitoring efficiency and provides higher frame-rates while being lightweight. The results are promising and as future work we will further explore what the network actually learns through its internal representations, and deploy it in a real experimental setup for further evaluation.

Acknowledgment

The author would like to acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.




Figure 8: Example of detection/predictions and estimation for number of targets. A common mistake by HoG-SVM and YOLO involves missing targets. In addition, other source of mistake relate to wrong localization with the bounding box. In both cases leads the camera to move in the wrong direction. Finally ACDCNet is able to correctly regress the number of targets and chooses to not move the camera in these frames as the targets are mostly in its FoV. (Best viewed in color)




Figure 9: Example of applying the approach to real-world images. The motion vector shows the predicted motion that the camera will make to keep the targets in its field of view. (Best viewed in color)

References

- [1] C. Micheloni, B. Rinner, and G. L. Foresti, "Video analysis in pan-tilt-zoom camera networks," *IEEE Signal Processing Magazine*, vol. 27, no. 5, pp. 78–90, Sept 2010.
- [2] C. Kyrkou, E. G. Christoforou, S. Timotheou, T. Theocharides, C. Panayiotou, and M. Polycarpou, "Optimizing the detection performance of smart camera networks through a probabilistic image-based model," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 28, no. 5, pp. 1197–1211, May 2018.
- [3] R. Wang, H. Dong, T. X. Han, and L. Mei, "Robust tracking via monocular active vision for an intelligent teaching system," *The Visual Computer*, vol. 32, no. 11, pp. 1379–1394, Nov 2016. [Online]. Available: <https://doi.org/10.1007/s00371-015-1206-8>
- [4] H. Chen, X. Zhao, and M. Tan, "A novel pan-tilt camera control approach for visual tracking," in *Proceeding of the 11th World Congress on Intelligent Control and Automation*, June 2014, pp. 2860–2865.
- [5] "Detection and tracking of moving objects; a survey," *International Journal of Engineering Research and Applications*, vol. 5, no. 11, 2015.
- [6] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Detect to track and track to detect," in *International Conference on Computer Vision (ICCV)*.

- [7] B. Bhanu, C. V. Ravishankar, A. K. Roy-Chowdhury, H. Aghajan, and D. Terzopoulos, *Distributed Video Sensor Networks*, 1st ed. Springer Publishing Company, Incorporated, 2011.
- [8] M. Al Haj, C. Fernández, Z. Xiong, I. Huerta, J. González, and X. Roca, *Beyond the Static Camera: Issues and Trends in Active Vision*. London: Springer London, 2011, pp. 11–30.
- [9] C. Ding, J. H. Bappy, J. A. Farrell, and A. K. Roy-Chowdhury, “Opportunistic image acquisition of individual and group activities in a distributed camera network,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 27, no. 3, pp. 664–672, March 2017.
- [10] S.-N. Lim, A. Elgammal, and L. S. Davis, “Image-based pan-tilt camera control in a multi-camera surveillance environment,” in *Multimedia and Expo, 2003. ICME '03. Proceedings. 2003 International Conference on*, vol. 1, July 2003, pp. I–645–8 vol.1.
- [11] M. A. Haj, A. D. Bagdanov, J. Gonzalez, and F. . Roca, “Reactive object tracking with a single ptz camera,” in *2010 20th International Conference on Pattern Recognition*, Aug 2010, pp. 1690–1693.
- [12] W. Luo, P. Sun, Y. Mu, and W. Liu, “End-to-end active object tracking via reinforcement learning,” 05 2017.
- [13] J. Ferryman and A. Shahrokni, “Pets2009: Dataset and challenge,” in *2009 Twelfth IEEE International Workshop on Performance Evaluation of Tracking and Surveillance*, Dec 2009, pp. 1–6.
- [14] F. Chollet, “keras,” <https://github.com/fchollet/keras>, 2015.
- [15] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, “Tensorflow: A system for large-scale machine learning,” in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI’16. Berkeley, CA, USA: USENIX Association, 2016, pp. 265–283. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3026877.3026899>
- [16] J. Redmon and A. Farhadi, “Yolo9000: Better, faster, stronger,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 6517–6525.
- [17] N. Dalal and B. Triggs, “Histograms of oriented gradients for human detection,” in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, vol. 1, June 2005, pp. 886–893 vol. 1.
- [18] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.