EdgeMA: Model Adaptation System for Real-Time Video Analytics on Edge Devices

Liang Wang[1,*], Nan Zhang[2,*], Xiaoyang Qu[2,✉], Jianzong Wang[2], Jiguang Wan[1],
Guokuan Li[1], Kaiyu Hu[3], Guilin Jiang[4], Jing Xiao[2]

[1] Huazhong University of Science and Technology, Wuhan, China
[2] Ping An Technology (Shenzhen) Co., Ltd., Shenzhen, China
[3] Stony Brook University, Stony Brook, NY
[4] Hunan Chasing Financial Holdings Co., Ltd., Changsha, China

Abstract. Real-time video analytics on edge devices for changing scenes remains a difficult task. As edge devices are usually resource-constrained, edge deep neural networks (DNNs) have fewer weights and shallower architectures than general DNNs. As a result, they only perform well in limited scenarios and are sensitive to data drift. In this paper, we introduce EdgeMA, a practical and efficient video analytics system designed to adapt models to shifts in real-world video streams over time, addressing the data drift problem. EdgeMA extracts the gray level co-occurrence matrix based statistical texture feature and uses the Random Forest classifier to detect the domain shift. Moreover, we have incorporated a method of model adaptation based on importance weighting, specifically designed to update models to cope with the label distribution shift. Through rigorous evaluation of EdgeMA on a real-world dataset, our results illustrate that EdgeMA significantly improves inference accuracy.

Keywords: Edge Computing · Deep Neural Network · Video Analytics · Data Drift · Model Adaptation.

1 Introduction

Real-time video analytics has become a promising application in the field of computer vision, which is powered by deep neural network (DNN) models, e.g., ResNet [7] and EfficientNet [28]. Video analytics applications such as traffic monitoring [17] use local cameras that continuously generate high-quality video streams to understand the environment. Most of these applications have to be carried out with real-time feedback. Therefore, edge computing is favored in video analytics because it eliminates the need for costly network overhead associated with uploading videos to the cloud and also reduces latency.

Due to limited computing resources, edge devices typically use lightweight DNN models [8, 34] for video analytics. Lightweight models have fewer weights

---

* These two authors have contributed to this work equally.

✉ Corresponding author: Xiaoyang Qu (e-mail: quxiaoy@gmail.com).

(a) Video snapshot (day)

(b) Video snapshot (night)

(c) Domain shift
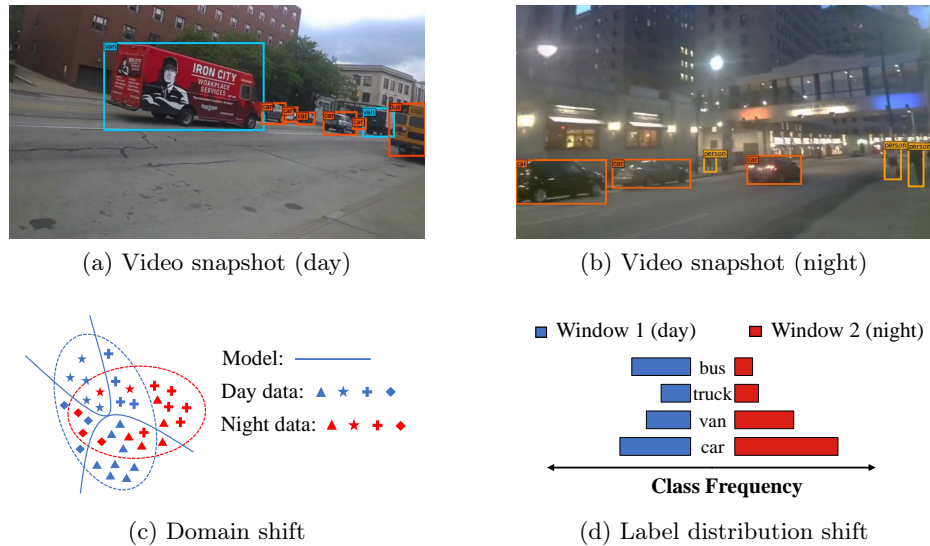
(d) Label distribution shift

Fig. 1. (a) and (b) show the crossroad traffic in the video during the day and night, respectively. (c) shows the domains from day to night and the results of inference labels for data. We count the frequency of objects within each time window, and (d) shows the statistics results of class label distribution.

and shallower architectures, often making them unsuitable for delivering high accuracy. In addition, when employed in the actual field, they are vulnerable to the data drift [21] problem: a mismatch between initial training data and live video data. Data drift results in substantial declines in accuracy as it violates the underlying assumption of DNN models that the training data from the past should resemble the test data in the future. This further aggravates the challenge of real-time video analytics.

In video analytics serving, domain and label distribution shifts are typical data drift phenomena [1, 12, 15, 16, 19]. As shown in Figure 1, street cameras encounter varying scenes over time. The label distribution of video varies over time, reducing the edge model's accuracy. In addition, the domain goes from day to night. While the model trained on the data collected based on daytime conditions fails to work well when deployed in the dark, degrading the accuracy.

Having a static model throughout the entire life-long inference process often negatively impacts performance, particularly with edge models. Due to their constrained capacity to learn object variations, edge DNNs require regular updates to accommodate changing object distributions and ensure optimal accuracy. We must consider the migration from domain to domain (e.g., weather, light).

In this work, we propose EdgeMA (**Edge M**odel **A**dapter), a video analytics system that resolves data drift in real-time video analytics on edge devices. EdgeMA copes with changes in the image domain and the label distribution

over time, providing the model adaptation for edge computing. In summary, the main contributions of our work are the following:

- We propose a novel system EdgeMA for real-time video analytics on edge devices to cope with both domain shift and label distribution shift.
- We design the lightweight domain detector to effectively identify and address domain shift.
- We detect label distribution shift and allow for real-time model retraining while ensuring the overall system remains resource-efficient.
- In our experiment on a real-world dataset, EdgeMA substantially outperforms scenarios without automatic model adaptation.

2 Related Work

Video analytics systems. Existing work has contributed to the creation of efficient video analytics systems. VideoStorm [33] investigates quality-lag requirements in video queries. Focus [9] uses low-cost models to index videos. Chameleon [13] exploits correlations in camera content to amortize profiling costs. NoScope [14] enhances inference speed through cascading models and filtering mechanisms. Clownfish [23] amalgamates inference outcomes from optimized models on edge devices with delayed results from comprehensive models in the cloud. However, these systems predominantly employ offline-trained models, overlooking the potential impact of model adaptation and data drift on accuracy.

Edge computing. The proliferation of the Internet of Things (IoT) has driven the production and usage of diverse hardware devices/sensors across the globe. These devices are able to collect data and then send it to the server for storage or processing, allowing end-users to access and extract the information as per their requirements [27]. Nevertheless, cloud computing has begun to reveal some issues. The centralized nature of cloud computing, handling data generated by global end devices, leads to numerous challenges, such as reduced throughput, increased latency, bandwidth constraints, data privacy concerns, and augmented costs. These challenges become particularly pressing in IoT applications, which demand rapid and low-latency data processing, analytics, and result delivery [24, 25]. To combat the aforementioned challenges associated with cloud computing, a novel computing paradigm known as edge computing, has attracted widespread attention [10]. In essence, edge computing offloads data processing, storage, and computing tasks traditionally assigned to the cloud to the network's edge, in close proximity to the terminal devices. This transition helps minimize data transmission and device response times, alleviate network bandwidth pressure, decrease data transmission costs, and facilitate a decentralized system [5].

Data drift. Edge DNNs can only memorize a limited number of object appearances and scenarios. Hence, they are particularly vulnerable to data drift [12,21], which arises when real-time video data significantly diverges across domains. Variations of scene density (for instance, during rush hour) and lighting conditions (such as daytime versus nighttime) over time pose challenges for surveillance cameras attempting precise object detection. Additionally, the distribution

of object classes changes over time, leading to a decrease in the precision of the edge model [29]. Because of their restricted ability to memorize changes, edge models necessitate continuous retraining with the most recent data and shifting object distributions to maintain high accuracy. Video-analytics systems are beginning to adopt continuous learning to adapt to changing video scenes and improve inference accuracy [1, 15, 16]. Ekya [1] provides sophisticated resource-sharing mechanisms for efficient model retraining. AMS [16] dynamically adjusts the frame sampling rate on edge devices depending on scene changes, mitigating the need for frequent retraining. RECL [15] combines Ekya and AMS, and offers more rapid responses by choosing an appropriate model from a repository of historical models during analysis. EdgeMA further considers lightweight retraining processing at the edge.
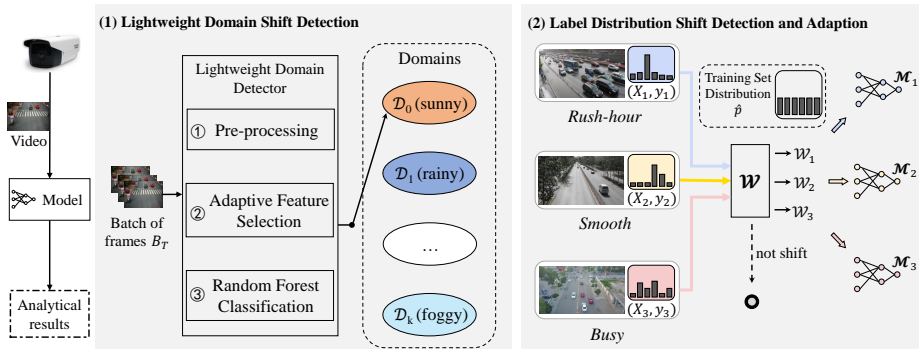
3   System Design

Fig. 2. EdgeMA system overview.

Figure 2 presents an overview of EdgeMA. Each edge device continuously performs inference for video analytics and buffers sampled video frames over time. The EdgeMA at the edge device runs iteratively on each new buffered batch of frames B_T, deciding whether to adapt the lightweight specialized model. The fundamental information upon which our system depends, when analyzing video feeds from static cameras or any other videos where the camera remains stationary, is that their domains and label distributions exhibit spatio-temporal locality. The model adaptation algorithm consists of two phases: (1) lightweight domain shift detection and (2) label distribution shift detection and adaptation.

In the first phase, EdgeMA detects in which domain the current batch lies. In our system, each domain corresponds to a training set containing a few data samples (e.g., 2000 images) from the domain scenario and an existing model pre-trained on a general dataset. After detecting, if the domain has changed,

the model is fine-tuned using the corresponding training set to adapt to the specialized domain.

In the second phase, the edge device detects if the label distribution shifts. If the label distribution changes, the model also requires retraining with the dataset of the current domain. We assume that it makes sense to give more importance to the classes of training data whose frequency of appearance is nearest to the features of the live data to improve accuracy for different label distributions.

When our training set is small, retraining can perform well with low computing costs [6]. And for concreteness, we describe our design for object detection, but the system is general and can be adapted to other tasks. Next, we illustrate the two phases of model adaptation in detail.

3.1 Lightweight Domain Shift Detection

The variations of illumination, weather, or style in different domains can seriously affect the model performance. Cross-domain detection for video is the challenge that model adaptation must overcome. Previous works [11, 32] can use heavily computation-intensive networks for domain shift detection. However, edge devices must allocate computing resources as much as possible to the video analytics model. It is not desirable to run the video analytics model with the complex domain detection network simultaneously on a resource-constrained edge device. Therefore, we design a lightweight domain detector for EdgeMA. We detect the video's environment condition by a random forest classifier [3] (RFC) and Gray Level Co-occurrence Matrix (GLCM)-based texture features, to detect the video's environmental conditions. This approach not only ensures high accuracy but also minimizes computational expenses and conserves energy.

Pre-processing Operation. EdgeMA requires conversion of the frame to grayscale, wherein each pixel value represents spectral intensity ranging between 0 and 255. The grayscale image, represented by g, is derived by computing the luminance of the color image using equation (1):

$$g = 0.299 \times R + 0.587 \times G + 0.114 \times B \qquad (1)$$

Adaptive Feature Selection. For image classification, texture features, dependent on the repetitive patterns across an image region, are predominantly utilized. The GLCM of an image I of dimensions $n * m$ characterizes the texture by quantifying occurrences of pixels with specific absolute values based on a spatial offset, is defined as:

$$C(g_i, g_j) = \sum_{p=1}^{n} \sum_{q=1}^{m} \begin{cases} 1, \text{if } I(p,q) = g_i \text{ and } I(p + \Delta_x, q + \Delta_y) = g_j \\ 0, \text{otherwise} \end{cases} \qquad (2)$$

where g_i and g_j represent the gray level values of an image. For spatial positions p and q within the image I, the offset (Δ_x, Δ_y) is determined by both the direction and distance for which the matrix is computed. GLCM features are computed with directions $\theta \in \{0°, 45°, 90°, 135°\}$, distance $d \in [1, 30]$, and the 6

prevailing texture properties [20], including contrast, correlation, homogeneity, angular second moment, dissimilarity, and energy. In total, 4*30*6=720 features are derived from GLCM.

Moreover, we employ the AdaBoost method [26] for feature selection to ascertain the relative importance of features. AdaBoost not only selects the most significant features but also assigns weights to weak classifiers to enhance classification performance. The algorithm based on AdaBoost starts by choosing a feature as an initial weak learner. During each iteration t, the primary steps of the selection process are shown as follows. (1) Normalize the weights of the images as $w^{[t]} = \left\{ w_I^{[t]} \right\}_I$, where I is from the training image set. (2) Choose the feature F_j that yields the least classification error summed over all images, weighted by $w^{[t]}$. (3) Increase the value S_j, denoting the importance score of the feature F_j. (4) Adjust the weights for each image sample $I(w^{[t]})$. The weights should correlate with the error rate E for that image. Through this method, images correctly classified by the chosen feature have diminished influence, whereas the weights of misclassified images increase. The final output of the algorithm is the vector S, signifying the relative importance of the original features.

Random Forest Classification. RFC [3] is a prominent method in machine learning, particularly for high-dimensional classification. RFC comprises a collection of decision trees, with each tree generated using a random vector drawn independently from the feature vector. During training, each tree's training set is formed from a bootstrap sample of the data, selecting frames with replacement. Each tree then casts a vote for the most probable class to predict the output. To develop and evaluate our proposed model, the classifier was trained multiple times, recording the best accuracy achieved. The frames were classified into diverse domain classes.

When EdgeMA is executed, it, by default, retrieves the last ten frames from the batch, converts them to grayscale, and extracts features to ascertain the domain. To evaluate our proposed detector, we train it multiple times and record the highest accuracy, as discussed in Section 4.1.

3.2   Label Distribution Shift Detection and Model Adaptation

After identifying the specific domain, EdgeMA fine-tunes the model using the corresponding domain training set. Fine-tuning employs importance weighting [4] (IW). This powerful technique applies a weight of importance to each class in the training set based on the distributional feature, which captures similarity to the live data distribution.

The label distribution of the training set X_S for the current domain is $P_S(X_S)$ (default IID). And the label distribution of the dataset X_T consisting of all N frames in the batch B_T is $P_T(X_T)$. The distribution $P_T(X_T)$ of different windows is non-IID and dynamically varying. Figure 3 shows the streaming schema of main operations.

Estimate P_T. The model M identifies K classes in total, and the predicted value of these is $f(*)$, then the distribution $P_T(X_T)$ is as follows:

Fig. 3. Streaming schema of the main operations.

$$P_T(X_T) = \left\{ \sum_{j=1}^{N} \frac{\{f(x_j) == i\}}{N} \right\}_{i=1}^{K} \tag{3}$$

where $\{f(x_j) == i\}$ is a Boolean operation (1 if equal, 0 if not). The distribution outcome is a vector.

Compute W. Similarly to [18], the importance weight $W(X)$ is calculated as follows. It explains that the class is given a higher weight with a higher probability of occurrence:

$$W(X) = \frac{P_T(X_T)}{P_S(X_S)} = C_{h_0}^{-1} q_{h_0} \tag{4}$$

where h_0 is the initial neural network trained on X_S, C_{h_0} is the confusion matrix of h_0 in source label distribution P_S, and q_{h_0} is a K-dimension vector $q_{h_0}[i] = P_T(h_0(X) = y_i)$.

Fine-tune M. Finally, model M is fine-tuned with W to cope with the ever-changing distribution. To reduce time cost, we utilize *coordinate descent* [22,31]. In this approach, EdgeMA retrains a limited subset of parameters (e.g., 20%) during each fine-tuning phase, with the edge device refining them over k iterations. The final optimization target for fine-tuning the model is shown below:

$$\frac{1}{n} \sum_{i=1}^{n} \frac{P_T(y_i)}{P_S(y_i)} f(x_i, y_i) \rightarrow \mathbb{E}_{x,y \sim P_S} \left(\frac{P_T(y)}{P_S(y)} f(x, y) \right) \tag{5}$$

Notably, fine-tuning can take up computing resources on the edge device. EdgeMA needs to determine the malignancy of a shift and reduces model update frequency. EdgeMA allows the model lags *if not malignant*. In practice, distributions shift constantly, and often these changes are benign. We employ the Kullback-Leibler (KL) divergence to quantify the discrepancy between the distribution P_M at the final stage of fine-tuning and the current distribution P_T. The KL divergence is defined as:

$$d = D_{KL}(P_T \| P_M) = \mathbb{E}_{x \sim P_T} \left[\log \frac{P_T(X_T)}{P_M(X_M)} \right] \tag{6}$$

By configuring a distance threshold D to determine whether adaptive learning is required, if $d < D$, it means that the model still matches the current label distribution.

4   Evaluation

In this section, we assess the performance of EdgeMA. Specifically, our evaluation seeks to answer two primary questions: (1) How effective is the lightweight detector, and what constitutes its peak performance? (2) How does EdgeMA cope with domain and label distribution shifts, and how does the effect of model adaptation perform in comparison to the baseline model?

4.1   Effectiveness of Lightweight Domain Detector

We evaluate EdgeMA on the task of object detection with the UA-DETRAC [30] dataset, which contains various environments, including sunny, rainy, cloudy, and night. We first trained our lightweight detector multiple times to develop and assess it for detecting the four domain environments.


**Fig. 4.** Performance on different angles and distances.

**Fig. 5.** Feature Importance.

**Fig. 6.** Performance on number of best features.

Fig. 7. Classifier accuracy.

Figure 4 shows the accuracy for a collection of 6 features from GLCM using different pixel distances $d \in [1, 30]$ and different angles 0°, 45°, 90°, 135°. Then we select the 2 best neighboring pixel distances for each of the 4 angles (d={5,9} for 0°, d={4,11} for 45°, d={2,4} for 90°, d={2,6} for 135°). Consequently, the total number of features is calculated as 6*2*4 = 48 features.

The AdaBoost selects the top-performing features from the pool of 48 based on their feature importance. We set the iteration parameter for AdaBoost to 100. We number these 48 features. Among these, F_{22} has the highest score, which denotes the correlation of the GLCM matrix with d=4 and angle 45°. Figure 5 shows the importance of all features.

Then we analyze the accuracy by varying the number of top-ranked features based on their scores, as illustrated in Figure 6. The highest accuracy, 83.38%, is achieved with six selected features.

We set 32 as the number of decision trees in the above evaluation. In addition, we evaluate other numbers, and the results are shown in Figure 7. Assigning 32 trees for the random forest still shows the best performance.

4.2 Overall improvements of Model Adaptation

The UA-DETRAC dataset comprises over 140 thousand frames with timestamps, and we utilize these frames, concatenating part of them to form a one-hour-long video stream, 25 frames per second. We use Nvidia Jetson AGX Xavier boards as edge devices. For these devices, we use the YOLOv4 [2] with the Resnet18 [7] model backbone. We select an additional 2000 images as the training set, and we find that choosing the fine-tuning iteration k=8 is the ideal trade-off between accuracy and training time.


Fig. 8. Overall mAP.

Fig. 9. Top-1 Acc over time.

To validate the robustness and efficacy of our model adaptation, we compare the detection accuracy (mAP metric) of the static lightweight YOLOv4 and the models fine-tuned in different domain conditions. The results are shown in Figure 8. The adaptative solution delivers higher detection accuracy since the model is specifically adapted to the revealed shift. Figure 9 shows the result of the lightweight model and model adaptation on the inference top-1 accuracy of the video over time. There are several intervals when the accuracy significantly drops (about 3-5 minutes per interval) as the model lags during scene changes. Nevertheless, the average performance of model adaptation is much higher than using only a lightweight model for inference. The adaptive solution has a top-1

accuracy above 0.8 most of the time, while the fixed lightweight model solution is stable in the range of 0.6-0.7.

Table 1. Comparison of two different schemes.

Scheme	Fraction of Parameters	mAP@0.5	FPS$_{avg}$	Fine-tuning Time Cost
Lightweight	-	0.406	59.6	-
Model Adaptation	5%	0.470	56.3	95 s
	10%	0.493	54.9	123 s
	20%	0.515	52.2	192 s
	100%	0.523	37.5	887 s

Table 1 summarizes the performance of two different schemes. Fractions for model adaptation represent the percentage of how many model parameters for each fine-tuning. Selecting only 20% of model parameters performs very effectively, achieving >0.1 better mAP score than the lightweight model. It results in only a 0.008 loss of mAP on average, but it reduces the fine-tuning time cost from 887s for full-model fine-tuning to 192s. In addition, due to the usage of computing resources for model fine-tuning, it shows a slight speed advantage over the model adaptation scheme.

5   Conclusion

In this paper, we propose EdgeMA, an innovative framework designed for video analytics on edge devices, which utilizes model adaptation supported by data drift detection and adaptive retraining methods. EdgeMA fine-tunes the model based on importance weighting after detecting shifts in the domain and label distribution during inference. Our evaluation shows that EdgeMA delivers overall higher accuracy compared to the static setting without model adaptation.

References

1. Bhardwaj, R., Xia, Z., Ananthanarayanan, G., Jiang, J., Shu, Y., Karianakis, N., Hsieh, K., Bahl, P., Stoica, I.: Ekya: Continuous learning of video analytics models on edge compute servers. In: 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22). pp. 119–135 (2022)
2. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934 (2020)

3. Cutler, A., Cutler, D.R., Stevens, J.R.: Random forests. In: Ensemble machine learning, pp. 157–175. Springer (2012)

4. Fang, T., Lu, N., Niu, G., Sugiyama, M.: Rethinking importance weighting for deep learning under distribution shift. Advances in Neural Information Processing Systems **33**, 11996–12007 (2020)

5. Ghosh, A.M., Grolinger, K.: Edge-cloud computing for internet of things data analytics: Embedding intelligence in the edge with deep learning. IEEE Transactions on Industrial Informatics **17**(3), 2191–2200 (2020)

6. Gu, Y., Du, Z., Zhang, H., Zhang, X.: An efficient joint training framework for robust small-footprint keyword spotting. In: Neural Information Processing: 27th International Conference, ICONIP 2020, Bangkok, Thailand, November 23–27, 2020, Proceedings, Part I. pp. 12–23. Springer (2020)

7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)

8. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)

9. Hsieh, K., Ananthanarayanan, G., Bodik, P., Venkataraman, S., Bahl, P., Philipose, M., Gibbons, P.B., Mutlu, O.: Focus: Querying large video datasets with low latency and low cost. In: 13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18). pp. 269–286 (2018)

10. Hua, H., Li, Y., Wang, T., Dong, N., Li, W., Cao, J.: Edge computing with artificial intelligence: A machine learning perspective. ACM Computing Surveys **55**(9), 1–35 (2023)

11. Ibrahim, M.R., Haworth, J., Cheng, T.: Weathernet: Recognising weather and visual conditions from street-level images using deep residual learning. ISPRS International Journal of Geo-Information **8**(12), 549 (2019)

12. Jia, Y., Zhang, X., Lan, L., Luo, Z.: Counterfactual causal adversarial networks for domain adaptation. In: Neural Information Processing: 29th International Conference, ICONIP 2022, Virtual Event, November 22–26, 2022, Proceedings, Part VI. pp. 698–709. Springer (2023)

13. Jiang, J., Ananthanarayanan, G., Bodik, P., Sen, S., Stoica, I.: Chameleon: scalable adaptation of video analytics. In: Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication. pp. 253–266 (2018)

14. Kang, D., Emmons, J., Abuzaid, F., Bailis, P., Zaharia, M.: Noscope: Optimizing neural network queries over video at scale. Proceedings of the VLDB Endowment **10**(11) (2017)

15. Khani, M., Ananthanarayanan, G., Hsieh, K., Jiang, J., Netravali, R., Shu, Y., Alizadeh, M., Bahl, V.: Recl: Responsive resource-efficient continuous learning for video analytics. In: 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23). pp. 917–932 (2023)

16. Khani, M., Hamadanian, P., Nasr-Esfahany, A., Alizadeh, M.: Real-time video inference on edge devices via adaptive model streaming. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4572–4582 (2021)

17. Liao, X.C., Qiu, W.J., Wei, F.F., Chen, W.N.: Combining traffic assignment and traffic signal control for online traffic flow optimization. In: Neural Information Processing: 29th International Conference, ICONIP 2022, Virtual Event, November 22–26, 2022, Proceedings, Part VI. pp. 150–163. Springer (2023)

18. Lipton, Z., Wang, Y.X., Smola, A.: Detecting and correcting for label shift with black box predictors. In: International conference on machine learning. pp. 3122–3130. PMLR (2018)
19. Liu, C., Qu, X., Wang, J., Xiao, J.: Fedet: A communication-efficient federated class-incremental learning framework based on enhanced transformer. arXiv preprint arXiv:2306.15347 (2023)
20. Maji, K., Sharma, R., Verma, S., Goel, T.: Rvfl classifier based ensemble deep learning for early diagnosis of alzheimer's disease. In: Neural Information Processing: 29th International Conference, ICONIP 2022, Virtual Event, November 22–26, 2022, Proceedings, Part III. pp. 616–626. Springer (2023)
21. Moreno-Torres, J.G., Raeder, T., Alaiz-Rodríguez, R., Chawla, N.V., Herrera, F.: A unifying view on dataset shift in classification. Pattern recognition **45**(1), 521–530 (2012)
22. Nesterov, Y.: Efficiency of coordinate descent methods on huge-scale optimization problems. SIAM Journal on Optimization **22**(2), 341–362 (2012)
23. Nigade, V., Wang, L., Bal, H.: Clownfish: Edge and cloud symbiosis for video stream analytics. In: 2020 IEEE/ACM Symposium on Edge Computing (SEC). pp. 55–69. IEEE (2020)
24. Qin, M., Chen, L., Zhao, N., Chen, Y., Yu, F.R., Wei, G.: Power-constrained edge computing with maximum processing capacity for iot networks. IEEE Internet of Things Journal **6**(3), 4330–4343 (2018)
25. Qu, X., Wang, J., Xiao, J.: Quantization and knowledge distillation for efficient federated learning on edge devices. In: 2020 IEEE 22nd International Conference on High Performance Computing and Communications; IEEE 18th International Conference on Smart City; IEEE 6th International Conference on Data Science and Systems (HPCC/SmartCity/DSS). pp. 967–972. IEEE (2020)
26. Rojas, R., et al.: Adaboost and the super bowl of classifiers a tutorial introduction to adaptive boosting. Freie University, Berlin, Tech. Rep (2009)
27. Shi, W., Dustdar, S.: The promise of edge computing. Computer **49**(5), 78–81 (2016)
28. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
29. Wang, L., Lu, K., Zhang, N., Qu, X., Wang, J., Wan, J., Li, G., Xiao, J.: Shoggoth: Towards efficient edge-cloud collaborative real-time video inference via adaptive online learning. arXiv preprint arXiv:2306.15333 (2023)
30. Wen, L., Du, D., Cai, Z., Lei, Z., Chang, M., Qi, H., Lim, J., Yang, M., Lyu, S.: UA-DETRAC: A new benchmark and protocol for multi-object detection and tracking. Computer Vision and Image Understanding (2020)
31. Wright, S.J.: Coordinate descent algorithms. Mathematical Programming **151**(1), 3–34 (2015)
32. Xie, R., Yu, F., Wang, J., Wang, Y., Zhang, L.: Multi-level domain adaptive learning for cross-domain detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
33. Zhang, H., Ananthanarayanan, G., Bodik, P., Philipose, M., Bahl, P., Freedman, M.J.: Live video analytics at scale with approximation and delay-tolerance. In: 14th USENIX Symposium on Networked Systems Design and Implementation (2017)
34. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6848–6856 (2018)