

# EdgeFM: Leveraging Foundation Model for Open-set Learning on the Edge

Bufang Yang

The Chinese University of Hong Kong  
Hong Kong SAR, China  
bfyang@link.cuhk.edu.hk

Lixing He

The Chinese University of Hong Kong  
Hong Kong SAR, China  
1155170464@link.cuhk.edu.hk

Neiwen Ling

The Chinese University of Hong Kong  
Hong Kong SAR, China  
lingnw@link.cuhk.edu.hk

Zhenyu Yan<sup>†</sup>

The Chinese University of Hong Kong  
Hong Kong SAR, China  
zyyan@cuhk.edu.hk

Guoliang Xing

The Chinese University of Hong Kong  
Hong Kong SAR, China  
glxing@cuhk.edu.hk

Xian Shuai

Noah's Ark Lab, Huawei  
Technologies, Hong Kong SAR, China  
shuaixian1@huawei.com

Xiaozhe Ren

Noah's Ark Lab, Huawei  
Technologies, Hong Kong SAR, China  
renxiaozhe@huawei.com

Xin Jiang

Noah's Ark Lab, Huawei  
Technologies, Hong Kong SAR, China  
Jiang.Xin@huawei.com

## ABSTRACT

Deep Learning (DL) models have been widely deployed on IoT devices with the help of advancements in DL algorithms and chips. However, the limited resources of edge devices make these on-device DL models hard to be generalizable to diverse environments and tasks. Although the recently emerged foundation models (FMs) show impressive generalization power, how to effectively leverage the rich knowledge of FMs on resource-limited edge devices is still not explored. In this paper, we propose EdgeFM, a novel edge-cloud cooperative system with open-set recognition capability. EdgeFM selectively uploads unlabeled data to query the FM on the cloud and customizes the specific knowledge and architectures for edge models. Meanwhile, EdgeFM conducts dynamic model switching at run-time taking into account both data uncertainty and dynamic network variations, which ensures the accuracy always close to the original FM. We implement EdgeFM using two FMs on two edge platforms. We evaluate EdgeFM on three public datasets and two self-collected datasets. Results show that EdgeFM can reduce the end-to-end latency up to 3.2x and achieve 34.3% accuracy increase compared with the baseline.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Human-centered computing** → **Ubiquitous and mobile computing**; • **Computer systems organization** → **Embedded and cyber-physical systems**.

<sup>†</sup> Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*SenSys '23, November 12–17, 2023, Istanbul, Turkiye*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-0414-7/23/11...\$15.00  
<https://doi.org/10.1145/3625687.3625793>

## KEYWORDS

Foundation Models, Edge Computing, Offloading, Edge-cloud Collaborative System, Open-set Recognition, Internet of Things

## ACM Reference Format:

Bufang Yang, Lixing He, Neiwen Ling, Zhenyu Yan<sup>†</sup>, Guoliang Xing, Xian Shuai, Xiaozhe Ren, and Xin Jiang. 2023. EdgeFM: Leveraging Foundation Model for Open-set Learning on the Edge. In *The 21st ACM Conference on Embedded Networked Sensor Systems (SenSys '23), November 12–17, 2023, Istanbul, Turkiye*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3625687.3625793>

## 1 INTRODUCTION

Deep learning models have been widely deployed on IoT systems thanks to their excellent performance and the advancement in edge Artificial intelligence (AI) hardware. There will be more than 1.5 billion edge AI processors shipped in 2024 [13]. Various commercial and industrial applications are deployed on embedded AI systems, such as health monitoring systems [49, 70], service and logistics robots [53], and sound event detection systems [27, 67]. As these applications must operate in complicated and ever-changing environments, scalability and adaptiveness are of great importance.

Most current on-device AI models are task-specific and can only predict a *closed-set* of classes pre-defined at the training stage [49, 68, 71]. Their performance degrades severely when the class of input is not seen during the training. Although various approaches such as transfer learning [2, 42] and meta-learning [15, 19] have been proposed to calibrate models on the edge, they still require non-trivial efforts for manual data annotation and on-device training, which are not practical in real-world deployments [33, 44, 49]. The recent emergence of foundation models (FMs) such as GPT [5] and CLIP [52] have shown impressive general knowledge that can support diverse downstream tasks like image captioning, question answering, and information extraction. Current commercial proprietary FMs usually contain millions or billions of parameters and are pre-trained using billions of data samples [4, 37]. Moreover,

multi-modal FMs such as CLIP and ImageBind can learn embedded matching through massive paired data of different modalities, such as text, images, audio, and motion data. Their general knowledge can process diverse sensor data with **open-set recognition** capability when the interested types of classes change.

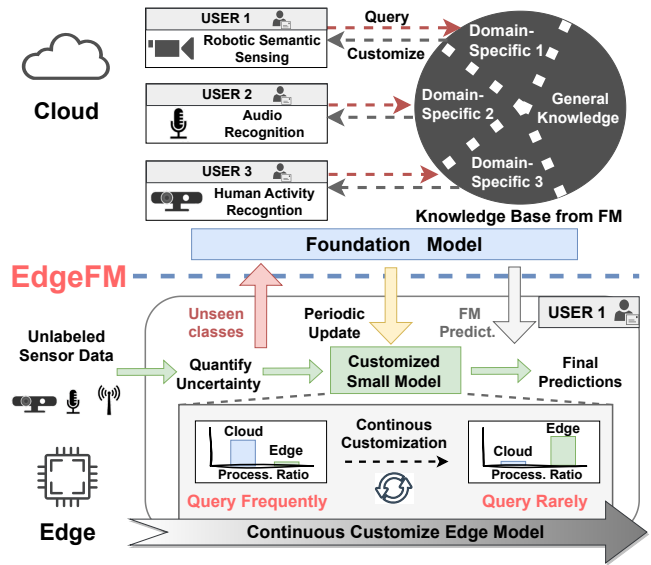
Some systems execute FMs on the cloud for tasks like robotic navigation [1, 38, 53]. However, transmitting all raw data to the cloud is not feasible in many practical scenarios. Executing bulky FMs on the edge directly is infeasible due to limited resources. Current model compression techniques [58, 65, 75] treat all samples equally during inference, which leads to significant performance degradation for difficult/unseen input data. Several studies partition a large model and deploy them to the edge and cloud for collaborative execution [29, 34, 76]. Most FMs adopt transformers whose output is even larger than the data input [63], so model partitioning of FMs is not desirable.

In this paper, we propose EdgeFM, a novel edge-cloud cooperative system that can achieve open-set recognition capability by leveraging FMs for selective knowledge query and edge model customization. As shown in Figure 1, FM is deployed on a cloud server and acts as a knowledge base containing both general and domain-specific knowledge. At the inference stage, EdgeFM first determines whether it should query FMs based on the uncertainty of semantic features of sensor data and the real-time network variations, which ensures the accuracy is always close to the original FM. Meanwhile, EdgeFM selectively uploads unlabeled data to query the FM on the cloud and periodically customizes the domain-specific knowledge and architectures for small models in a *label-free manner*. When the data distribution or interested class set changes, EdgeFM will query the knowledge from FMs frequently at the early stage, while it can primarily execute customized small models on edge devices at the late stage, thus reducing system overhead subsequently. EdgeFM supports different tasks and modalities so that different users can query their domain-specific knowledge from FM while only executing their customized small models on resource-constrained edge devices to save system overhead.

We extensively evaluate the performance of EdgeFM on two FMs (CLIP and ImageBind), and two edge devices (Nvidia Jetson Xavier and Nano), using three public datasets and two self-collected datasets. Our results show that EdgeFM reduces end-to-end latency up to 3.2x compared with existing on-device inference approaches. EdgeFM can also achieve up to 34.3% accuracy increase compared with the existing on-device open-set recognition approach.

The contributions of this paper are summarized as follows:

- We propose EdgeFM, the first edge-cloud cooperative system with open-set recognition capability leveraging FMs for selective knowledge query and dynamic edge model customization. The system can work with sensor data of different modalities.
- We design a semantic-driven customization approach that allows EdgeFM to customize the domain-specific knowledge and the architectures of mobile-friendly models in a *label-free manner*.
- We develop a dynamic model switching approach considering both the uncertainty of semantic features of sensor data and the real-time network fluctuation.



**Figure 1: An example of EdgeFM, enabling edge devices with open-set capability using dynamic customization and runtime model switching.**

- We implement EdgeFM on two FMs and deploy it to PC and two edge platforms. The evaluation includes three public datasets and two real-world datasets collected by ourselves about daily activity recognition and robot semantic recognition. EdgeFM can reduce the end-to-end latency up to 3.2x and achieve a 34.3% accuracy increase compared with the baseline.

## 2 BACKGROUND AND RELATED WORK

### 2.1 Multi-Modal Foundation Models

*Foundation models (FMs)* refer to a new class of large machine learning models that can extract valuable features to support diverse downstream tasks such as chatbot (e.g., GPT [5]) and image recognition (e.g., CLIP [52]). Multi-modal FMs represented by CLIP and ImageBind [18, 52] learn the pairing of data of different modalities (e.g., RGB image, depth image, and audio) and their corresponding textual description across the internet to achieve the capability of *open-set recognition*. In particular, multi-modal FMs adopt a transformer-based encoder to extract features from the raw data and convert them into *embedding*. FMs also use a text encoder to extract *text embedding* from the corresponding textual description of the raw data (e.g., images). The training of such multi-modal FMs often adopts contrastive learning to study the pairing between the data embedding and the corresponding text embedding to construct a *unified embedding space*. This is also the reason that multi-modal FMs like ImageBind can work with one or multi-modal sensor data. Specifically, for any class described in a natural language manner, FMs such as CLIP first convert the class name *CLS* into a text description through concatenating *CLS* with a pre-defined text template called prompt, such as “This is a photo of a {CLS}”. Then, the text description and the sensor data (such as images) are fed into the text encoder and image encoder of FMs to obtain respective

embedding. After computing the similarity score between the text embedding and other sensor data's embedding, FMs select the class with the highest score as the final prediction.

## 2.2 Related Works

**Open-set Recognition.** Open-set recognition aims to recognize any classes described in a natural language manner without fine-tuning. Existing approaches for open-set recognition can be classified into two categories: semantic-based and generative adversarial network (GAN)-based. Semantic-based approaches [8, 27, 57, 61, 67] build the connection between the embeddings of sensor data and the semantic embedding of natural language. GAN-based approaches [46] use a GAN-based semantic decoder to synthesize features for unseen classes and mix them up with the real features of seen classes for model training. However, their training data and model parameters are still limited compared to FMs, showing unsatisfactory performance of open-set recognition on embedded edge systems.

**Efficient On-device Inference.** Many studies are proposed to accelerate NN inference on edge devices, such as model compression [48, 66], early-exit [6, 35], input filtering [25, 74] and reusing [16, 69]. However, model compression, such as quantization [66], pruning [48], and knowledge distillation (KD) [21], are static acceleration approaches. They will also suffer severe performance degradation when compressing FMs into the scale of lightweight CNNs. Early exit [6, 35] can dynamically reduce the redundant computation of NNs. However, the high-dimensional embeddings and deep layers of FMs make the early-exit heads very heavyweight. Moreover, they still require executing the entire FM on the edge for hard samples, which can exceed the memory bound of edge devices like Nvidia Jetson Nano. Some work proposes to optimize computation efficiency by processing input data, including filtering [25, 74] and reusing [16, 69]. However, they can still not address the challenge of insufficient memory on edge devices due to the bulky size of FMs.

**FMs on the Edge.** Although FMs are more recently emerging, some approaches have been proposed to optimize the inference of FMs. MLC-LLM [60] leverages memory planning and quantization techniques to run LLMs on the phone. Tabi [64] and FrugalGPT [7] propose to cascade different sizes of models for acceleration. However, these approaches are tailored for generating text dialogues, which cannot work with multi-modal sensor data. A line of work including DIME-FM [58], FD-CLIP [65], VLKD [12], and MobileSAM [75] aims to compress multi-modal FMs by KD. However, most of the previous works focus on preserving the great open-set ability of FMs rather than generating task-specific small models, thus always requiring heavyweight transformer-based architectures [63], which is hard to be implemented on embedded systems. These approaches also require the dedicated design of small model architectures, limiting their practicality.

**Edge-cloud Collaboration.** Several works propose to adopt edge and cloud collaboration solutions to achieve the trade-off between accuracy and efficiency, Neurosurgeon [29] splits the NN to deploy the several layers on the edge while the remaining layers are on the cloud. SPINN [34] integrates the model splitting and early-exit to co-optimize multiple objectives, including accuracy and latency, by

adjusting the split point and early exit threshold. AgileNN [26] and DeepCOD [73] compress the size of transmitted intermediate features to improve the edge-cloud inference efficiency. There are also edge-cloud cooperative systems based on big-little model-switching [36, 50]. However, most of the previous edge-cloud cooperative systems focus on task-specific NNs, which work in a closed-set manner. The development of an open-set supported cooperative system remains unexplored.

In summary, existing works either focus on optimizing the on-device efficient inference for closed-set NNs, or compressing FMs through dedicated designs that are static and not scalable for dynamic real-world IoT applications. Leveraging open-set knowledge of FMs for embedded systems is still an ill-address problem.

## 3 MOTIVATION: A CASE STUDY


The capability of open-set recognition is highly desirable in both commercial and industrial embedded systems. Current FM services predominantly adopt cloud-centric solutions [11, 39]. We first evaluate the performance of cloud-centric FM services under real-world dynamic network conditions. Next, we measure the performance of FMs and small-size recognition models on objects of unseen classes. This applies to many embedded systems installing an AI model that can only recognize a limited set of objects due to constrained resources. We use a public image dataset FLO102 [47] with 102 types of flowers, and an image dataset collected by us containing 40 classes of common objects in the indoor environments. We use CLIP [52] and ImageBind [18] as FMs and use lightweight models for embedded systems like MobileNetV2 [23] and ResNet18 [20] as small models. In particular, we first study the performance of FM and small models on unseen classes, the customization of small models to adapt to the dynamic set of classes, and the execution efficiency of FMs and small models. Then, we test the feasibility of customizing embedded models with FM's knowledge for the open-set capability.

### 3.1 Cloud-centric FM services

We first measure how the cloud-centric FM service performs on an edge system. We set up an edge platform (i.e., NVIDIA Jetson Nano) to stream RGB images of an office room to the cloud server for recognizing the objects. The cloud server deploys ImageBind as the FM. Figure 2 shows the network bandwidth measurements and corresponding inference latency. The inference latency exhibits considerable fluctuations under dynamic network conditions, ranging from 200 to 630 ms. However, cloud-centric approaches necessitate streaming all data to the cloud server, causing non-trivial system latency (up to 630 ms) due to varying network conditions. The increased delays can significantly affect user experience, such as the risk of collisions of home service robots.

### 3.2 Understanding FMs and SMs


**Open-set Capability.** We split the dataset into training and test datasets with no overlapped classes to test the open-set capability of FMs and small models (SMs). Table 1 shows that FMs can achieve up to 77% mean accuracy without any fine-tuning or calibration, while SMs only show 1.5% accuracy on average, which is equivalent to random guessing. The significant performance gap shows that



**Figure 2: An example showing the inference latency of cloud-centric FM solutions under dynamic network conditions.**

**Table 1: The performance of small models (SMs) and foundation model (FMs) with unseen test samples on FLO102 and SC40 datasets, respectively.**

	Models	FLO102	SC40	Param.	FLOPS	Nano
SMs	MobileNet	1.1%	2.6%	3.5M	0.3B	36.8ms
	ResNet18	0.4%	3.4%	11.7M	1.8B	30.5ms
FMs	ImageBind	78.4%	71.3%	1172M	167.3B	N.A.
	CLIP-L/14	79.5%	77.1%	407.8M	61.5B	N.A.



**Figure 3: The performance of SMs with different amounts of fine-tuned data. The parameters of FMs are frozen.**

**Figure 4: Feasibility of leveraging FMs as a knowledge base for customizing small models.**

FMs can be used for open-set capability while SMs are hardly usable when the test data is not seen during the training.

**Customization of Small Models.** We further test the performance of customized small models by calibrating them with data from the new classes (i.e., the new types of objects added to the environment). Figure 3 shows the accuracy of the small models customized by different amounts of labeled data from the new classes. SMs show unsatisfactory performance that is far below that of FMs when the amount of labeled data is limited. When more labeled data is used for calibration, SMs can achieve a good performance that is even higher than FMs, up to 92%. This shows that the fully customized small model can achieve similar and even superior recognition performance than FMs on specific tasks.

**Execution Latency of FMs and SMs.** Table 1 shows the parameter size and inference latency of customized SMs and FMs. The parameter size and computational overhead in FLOPS of MobileNetV2 are

335x and 558x less than ImageBind, respectively. The inference latency for one single image of MobileNetV2 and ResNet18 are 36.8ms and 30.5ms on Jetson Nano, respectively. However, both ImageBind and CLIP-L/14 require more than 6GB memory, which exceeds the memory limit of Jetson Nano, and thus can not directly execute (marked as N.A. in Table 1). The huge computation requirement of FMs urges it mostly deployed on the cloud.

The above preliminary results motivate us to adopt a model-switching solution between FMs on the cloud and customized small models on the edge, which can leverage the respective advantages of FMs and customized small models. In summary, FMs like CLIP and ImageBind have good open-set capabilities, but they are too large to fit in embedded systems. On the other side, SMs can achieve a better performance than FMs after customizing models with labeled data from the new classes. However, the labeled data for customization is usually not available in practical applications.

### 3.3 Customization with FMs

We further explore the feasibility of leveraging FMs to customize SMs by using the prediction results of FMs as pseudo labels to supervise the customization of SMs. Specifically, we train SMs under different percentages of correct labels and validate whether training with noisy labels is possible. Figure 4 shows the performance of SMs when we set different accuracy of the pseudo labels. We control the correct percentage of the pseudo labels among 400 unseen test samples, which is the values on the x-axis. The vertical line shows the accuracy of CLIP-L/14. The result shows that the accuracy of the SM is 90% when all the pseudo labels are correct, while the accuracy of the SM drops to 80% when the pseudo labels have 80% accuracy. Meanwhile, CLIP can provide accuracy with 79.5% without any fine-tuning. This motivates us to use FMs as a rich and general knowledge base to provide high-quality supervision for the customization of SMs.

## 4 APPLICATION AND SYSTEM OVERVIEW

EdgeFM is a novel edge-cloud cooperative system with open-set recognition ability by using FMs for selective knowledge query and edge model customization. Next, we will introduce the application scenario and overview of the system.

### 4.1 Application Scenarios

We use two examples to discuss the potential applications of EdgeFM. The first application is robot semantic sensing [44] which is widely used for home services [33] and industrial applications [62]. The main challenge for the current home-service robot system is the high diversity in personal items and dynamic environmental factors across households. In addition, the types of objects in the environment will change over time. For example, users sometimes buy or throw items. Current recognition models on home-service robots [33, 44] require labeled data for model retraining to adapt to the new classes of objects, which is not realistic and user-unfriendly. The second application is human activity recognition (HAR), which is an essential algorithm for many smart embedded systems for health, such as disease monitoring and fitness tracking. The main challenge for current HAR systems [49] is that there is a need for *model calibration*. During the first-time installation, the HAR system requires

the user to manually collect and label data to calibrate the model to the environment and the target activities of users. Moreover, it may require periodic calibration when the monitored activities change due to the changing health condition or the doctor’s advice.

## 4.2 System Architecture

EdgeFM is a novel edge-cloud cooperative system that enables edge devices with open-set recognition. Figure 5 shows the system overview of EdgeFM. In the customization stage, EdgeFM conducts *knowledge query and customization* (§ 5.1) and *dynamic edge update* (§ 5.2) to selectively upload unlabeled data to the cloud and customize the domain-specific knowledge and architectures for small models. Specifically, EdgeFM first conducts user device profiling on the edge to obtain information about the applications (e.g., tasks and modalities) and computation resources (e.g., memory constraints) of edge devices (§ 5.2.2). The profiling results are then employed by the model selection module to determine the appropriate architecture for small models on the edge (§ 5.1.2). Meanwhile, EdgeFM selectively uploads the unlabeled sensor data to the cloud (§ 5.2.1) for customization. EdgeFM will conduct semantic-driven customization on the cloud (§ 5.1.1) and periodically update the customized small model and text embedding pool to the edge. During the inference, EdgeFM inference engine employs the dynamic network adaptation module (§ 5.3.2) to continuously monitor the network condition and update the threshold searching table. Then, EdgeFM conducts dynamic model switching (§ 5.3.1) to determine whether query FM on the cloud or use customized small models for inference. The dynamic model switching policy of EdgeFM supports open-set models and also considers both the uncertainty of the sensor data and the dynamic network variation.

## 5 DESIGN OF EDGEFM

### 5.1 Knowledge Query and Customization

In this section, we will introduce how EdgeFM customizes the domain-specific knowledge and architectures for the small models in a label-free manner.

**5.1.1 Semantic-driven Customization.** The limited computation resources of edge platforms and real-time requirements of tasks usually require adopting mobile-friendly CNNs architectures rather than the heavyweight transformer [63]. The main challenge here is how to effectively customize the heterogeneous lightweight CNNs by the knowledge from FMs in a *label-free manner* while preserving open-set recognition capability. To address this challenge, we propose a semantic-driven customization approach. Next, we will introduce the components of our semantic-driven customization approach respectively.

**Heterogeneous Feature Mapping.** Unlike existing work distilling knowledge between similar architectures, EdgeFM conducts customization between heterogeneous models, i.e. FMs to mobile-friendly CNNs. Existing lightweight CNNs [23, 59] can be regarded as consisting of a convolution-based feature extractor and a task-specific classifier. However, FMs usually adopt transformer-based architectures [63], which encode the input images or spectrograms

into a sequence of tokens and extract context information. Multi-modal FMs further align the embeddings of vision or audio modality with the text embedding in a unified embedding space. The difference in embedding space and heterogeneous model architectures between FMs and small models makes the customization challenging. Therefore, we discard the task-specific classifier of the original small models and add a feature projection network on top of the original feature extractor of small models, which is defined as  $\mathbf{v}_i = \psi(\mathcal{S}(\mathbf{x}_i))$ , where  $\mathcal{S}(\cdot)$  denotes the feature extractor of the customized small model,  $\psi(\cdot)$  is the feature projection network. The architecture of the feature projection network  $\psi(\cdot)$  is a lightweight single-layer feed-forward network. It can ensure the features of the customized small model have the same dimension as FM’s unified embedding space, facilitating matching with text embeddings of FMs to enable open-set recognition. For example, the output features of MobileNetV2 have a dimension of 1280, while the unified embedding space for ImageBind is 1024. Consequently, the input and output dimensions of the feature projection network are 1280 and 1024, respectively.

**Knowledge Query from the Foundation Model.** It is impractical to obtain high-quality labeled data in embedded systems. Under this scenario, the direct way to customize domain-specific knowledge from FMs is to use Mean-Squared-Error (MSE) loss to pull closer the embedding of sensor modality (e.g., image) between FMs and small models. Figure 6 shows the recognition accuracy of the customized small model when fine-tuned with labeled data and when employing unlabeled data with MSE loss for knowledge distillation from FM. We can see that employing unlabeled data with MSE loss will lead to significant accuracy degradation compared with using labeled data for fine-tuning.

To address this challenge, we propose to leverage FMs to further customize the *user-specific knowledge* to the small models. Here we define the user-specific knowledge as the interested class set specified by users. Specifically, EdgeFM initially pre-stores a *text embedding pool* that contains text embeddings from a wide range of frequently-used classes. The text embeddings in the pool are computed by the text encoder of FMs on the cloud. In practical use, EdgeFM allows users to freely add their interested classes for respective applications. The text embeddings of these newly added classes are computed by FMs on the cloud and are added to the pool, which is then updated to the user device periodically (§ 5.2.2). Note that it does not mean requiring users to annotate each data, but only providing interested classes set.

Take vision tasks as an example, suppose the text embedding pool is  $\mathbf{T}$  and the visual encoder of FM is  $\mathcal{T}_v(\cdot)$ , FMs on the cloud first extract the visual embedding of the unlabeled data  $\mathbf{x}_i$  as  $\mathcal{T}_v(\mathbf{x}_i)$ . Then, EdgeFM will select a text embedding  $\mathbf{t}'_i$  from the text embedding pool with the highest similarity with  $\mathcal{T}_v(\mathbf{x}_i)$  as:

$$\mathbf{t}'_i = \operatorname{argmax}(\langle \mathcal{T}_v(\mathbf{x}_i), \mathbf{t}_k \rangle), \mathbf{t}_k \in \mathbf{T}, \quad (1)$$

where  $\mathbf{t}'_i$  is defined as the pseudo text embedding. We also assign  $\mathbf{t}'_i$  a confidence score as  $w_i = \langle \mathcal{T}_v(\mathbf{x}_i), \mathbf{t}'_i \rangle$ , i.e. the cosine similarity between  $\mathcal{T}_v(\mathbf{x}_i)$  and  $\mathbf{t}'_i$ . The obtained pseudo text embedding and its confidence score will be used for semantic-driven distillation.

**Semantic-driven Distillation Loss.** To further customize the user-specific knowledge, we propose semantic distillation loss. Here we take vision recognition as an example. Firstly, we adopt MSE loss



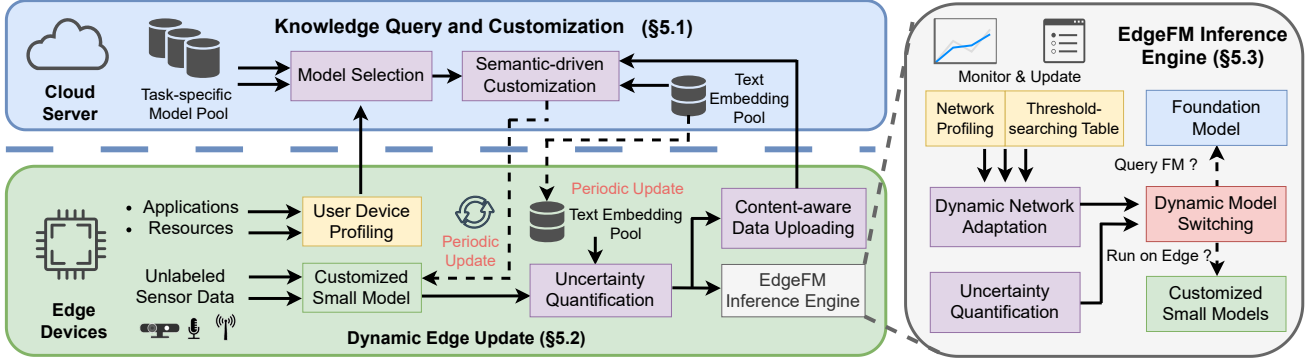


Figure 5: Overall system architecture of EdgeFM.

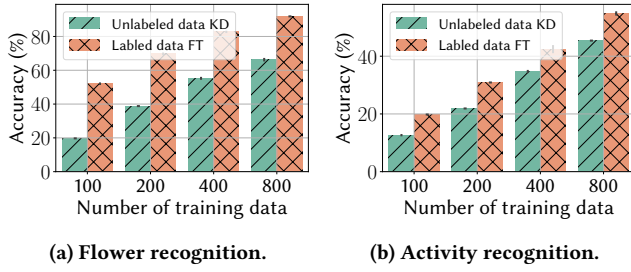


Figure 6: Customization accuracy with labeled data and unlabeled data. FT denotes fine-tuning.

to pull the features extracted by customized small models closer to the visual embedding of FMs, i.e.  $\mathcal{L}_{vis} = \mathcal{H}_{MSE}(\mathcal{T}_v(x_i), v_i)$ , where  $\mathcal{T}_v(x_i)$  and  $v_i$  are the visual embedding of FMs and customized small model, respectively;  $\mathcal{H}_{MSE}$  is MSE function. Next, we adopt a bidirectional contrastive learning loss [52] to further pull the features extracted by customized small models closer to the pseudo text embedding of FMs. Given a mini-batch of  $bs$  paired data  $(v_i, \hat{t}_i)$ , where  $v_i$  is the embedding of the sensor data (e.g., image) extracted by the customized small model, and  $\hat{t}_i$  is the most similar text embedding in the text embedding pool of FMs, as in Equation 1. We adopt the bidirectional contrastive learning loss which is defined as:

$$\mathcal{L}_i^{v \rightarrow t'} = -\log \frac{\exp\{\langle v_i, \hat{t}_k \rangle / \tau\}}{\sum_{k=1}^{bs} \exp\{\langle v_i, \hat{t}_k \rangle / \tau\}} \quad (2)$$

$$\mathcal{L}_i^{t' \rightarrow v} = -\log \frac{\exp\{\langle \hat{t}_i, v_k \rangle / \tau\}}{\sum_{k=1}^{bs} \exp\{\langle \hat{t}_i, v_k \rangle / \tau\}} \quad (3)$$

$$\mathcal{L}_{text} = \frac{1}{bs} \sum_{i=1}^{bs} w_i \left\{ \lambda \mathcal{L}_i^{v \rightarrow t'} + (1 - \lambda) \mathcal{L}_i^{t' \rightarrow v} \right\} \quad (4)$$

where  $w_i$  is the confidence score of the sample  $x_i$ , which is obtained as mentioned before.  $\lambda$  and  $\tau$  are the weight and temperature parameters, respectively. We test the two parameters with FLO102 dataset and choose  $\lambda = 0.5$  and  $\tau = 1$  that achieve the best recognition performance for the customized small model.

**5.1.2 Model Selection of Small Models.** We develop a model selection module that can customize the architecture of small models on the edge based on the tasks, modalities, and computation resources of edge devices. Figure 7 shows varied performance of four small models with different architectures on different tasks and data modalities. In particular, MobileNetV2 [22] performs better on vision-based recognition tasks like HAR, but has worse performance on the audio recognition task. This is because the depth-wise separable convolution is unsuitable for extracting the features from spectrogram-based data [23]. Moreover, the computation resources such as memory footprint and FLOPS vary for different edge devices and tasks. To this end, EdgeFM will determine the architecture of small models based on the tasks, data modalities, and computation resources of edge devices. Specifically, EdgeFM pre-stores many classical small model architectures with different accuracy and computation overhead such as MobileNet [23], and also those optimal architectures searched by Neural Architecture Search (NAS) technique such as EfficientNet series [59], on the cloud server. These small model architectures are grouped and stored in a task-specific model pool according to the modalities and tasks. At the offline stage, we test the recognition accuracy of each small model on public datasets and measure the resource usage, including the memory footprint and FLOPS. Note that we only use the model's accuracy on public datasets to assess its representation capability, without necessitating the use of labeled data from users. The accuracy, memory footprint, and FLOPS are recorded in a table, i.e., the accuracy-resource lookup table. At the online stage, EdgeFM will first select the corresponding model pool  $POOL_{app}$  based on the application specified by users. Next, EdgeFM determines the architecture of the small model by searching the accuracy-resource lookup table to maximize the recognition accuracy under the resource constraints of FLOPS and memory.

## 5.2 Dynamic Edge Update

**5.2.1 Content-aware Data Uploading.** Streaming all data to the cloud can cause non-trivial transmission overhead and unreliability. However, existing approaches [36, 50] utilize the softmax score by the conventional closed-set model to determine the data uploading, which is not suitable for open-set models. Therefore, we design a content-aware data-uploading approach tailored for open-set




Figure 7: Performance of small models with different architectures on different modalities.

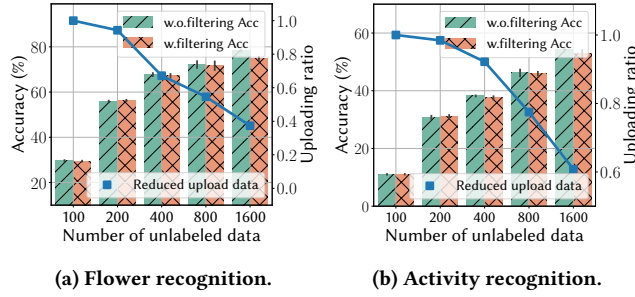


Figure 8: Customization accuracy and uploading data ratio with or without content-aware data-uploading.

models, which utilizes the semantic similarity between sensor data embeddings and text embeddings as uploading criterion.

Since our semantic-driven customization enables the open-set recognition capability of the customized small models, we use semantic similarity as uncertainty quantification of the samples. This distinguishes EdgeFM from previous studies in terms of the uncertainty metrics. Specifically, for each collected data samples  $\mathbf{x}_i$ , we first compute the cosine similarity between sensor data embedding  $\mathbf{v}_i$  (computed by the customized small models) and each text embedding  $\mathbf{t}_k$  in the text embedding pool  $\mathbf{T}$ , which is defined as  $\text{sim}(\mathbf{x}_i) = \langle \mathbf{v}_i, \mathbf{t}_k \rangle$ . We use the margin score [50] as uncertainty quantification for the samples, which is defined as  $\text{Unc}(\mathbf{x}_i) = \text{sim}_1(\mathbf{x}_i) - \text{sim}_2(\mathbf{x}_i)$ , where  $\text{sim}_1(\mathbf{x}_i)$  and  $\text{sim}_2(\mathbf{x}_i)$  are the highest similarity and the second highest similarity calculated between  $\mathbf{x}_i$  and all  $\mathbf{t}_k$  in the text embedding pool. Only samples with  $\text{Unc}(\mathbf{x}_i) < V_{thre}$  will be uploaded to the cloud for customization. We set  $V_{thre} = 0.99$  based on the observation that this configuration results in a substantial reduction in data transmission with negligible accuracy deterioration.

Figure 8 shows the customization accuracy and uploading data ratio with or without content-aware data uploading. We can see that with the collected unlabeled sensor data increasing from 100 to 1600, the ratio of uploading data decreases from 100% to about 40% on two applications (as the blue line shows) with a negligible accuracy drop. Therefore, content-aware data uploading can help EdgeFM reduce the network transmission overhead in real-world implementations.

**5.2.2 User Device Profiling and Periodic Update.** To reduce the transmission overhead, EdgeFM periodically updates the customized

small model and text embedding pool and delivers them to the edge device. Compared to prior studies [3, 30], a unique characteristic of our approach is the dynamic updating of the text embedding pool on the edge side, which enables the support of open-set recognition on edge devices. Specifically, EdgeFM utilizes a user device profiler to record the information of edge devices, such as applications (e.g., tasks and modalities), and computation resources of edge devices (e.g., memory usage and latency requirements). This information is then employed by the model selection module (§ 5.1.2) to select the appropriate architecture for the customized small model. On the other hand, EdgeFM continuously collects sensor data from the environment and selectively uploads them to the cloud server. Upon the uploaded data reaching the specified amount, EdgeFM will conduct semantic-driven customization and subsequently download the updated customized small model to the edge device. Moreover, the text embedding pool on the cloud is also updated if users add their interested class set. The text embedding pool and customized small model will be updated synchronously to the user device. The frequency of periodically updating the edge side in EdgeFM offers a trade-off between accuracy and transmission overhead. Since the experimental results in [3] have shown that setting the updating interval to 200 sec yields the best trade-off between accuracy and transmission overhead, EdgeFM adopts the same updating interval of 200 sec for both customized small models and the text embedding pool.

### 5.3 EdgeFM Inference Engine

This section introduces the EdgeFM inference engine, which performs dynamic model switching at runtime, considering both the uncertainty of sensor data and network variation.

**5.3.1 Dynamic Model Switching.** EdgeFM adopts an edge-cloud hybrid prediction mechanism based on the collaboration between customized small models, router model, and FM, where the first two models run on edge devices while FM runs on the cloud. The overall prediction of EdgeFM for input sample  $\mathbf{x}_i$  is defined as:

$$P(\hat{y} | \mathbf{x}_i) = r(\mathbf{x}_i)P_{SM}(\hat{y} | \mathbf{x}_i) + (1 - r(\mathbf{x}_i))P_{FM}(\hat{y} | \mathbf{x}_i) \quad (5)$$

where  $P_{SM}(\hat{y} | \mathbf{x}_i)$  and  $P_{FM}(\hat{y} | \mathbf{x}_i)$  are the predictions of the customized small model and FM, respectively. Note that the predictions of the open-set model in EdgeFM are computed by the cosine similarity score between sensor data embeddings and text embeddings, i.e.  $P_{SM}(\hat{y} | \mathbf{x}_i) = \langle \mathbf{v}_i, \mathbf{t}_k \rangle$ ,  $P_{FM}(\hat{y} | \mathbf{x}_i) = \langle \mathcal{T}_v(\mathbf{x}_i), \mathbf{t}_k \rangle$ .  $\mathbf{v}_i$  and  $\mathcal{T}_v(\mathbf{x}_i)$  are the sensor data embeddings that are computed by the customized small model and FM, respectively. A router model  $r(\mathbf{x}_i)$  controls the models switching according to the prediction of the customized small model and a threshold  $thre(t)$ , which is defined as:

$$r(\mathbf{x}_i) = \mathbb{1}\{\text{Unc}(\mathbf{x}_i) \geq thre(t)\} \quad (6)$$

where  $\text{Unc}(\mathbf{x}_i)$  is the uncertainty of the sensor data, which is defined in Section 5.2.1. Note that the threshold  $thre(t)$  in the inference engine is different from the threshold in the dynamic edge update. EdgeFM tunes the threshold  $thre(t)$  at runtime to adapt to the dynamic network condition (see Section 5.3.2). Based on the uncertainty of the input sample and threshold, the router model determines whether to query FMs on the cloud for inference or use the prediction of the customized small model.

**5.3.2 Dynamic Network Adaptation.** Model switching threshold determines the trade-off between accuracy and inference latency. EdgeFM adopts the dynamic network adaptation module to find the optimal model switching threshold under the dynamic network fluctuation.

Specifically, EdgeFM will periodically collect a specific number of sensor data from the environment to build a calibration set. We sample the threshold equally in the range of  $(0, 1)$ . For each  $thre \in (0, 1)$ , EdgeFM computes the edge-side processing proportion  $r(thre)$ , overall accuracy  $acc(thre)$ , edge-side processing latency  $t_{edge}$ , transmission latency  $t_{trans}$ , and cloud-side processing latency  $t_{cloud}$  for the calibration set, and saves them in a threshold-searching table. The estimated end-to-end inference latency can be expressed as:

$$\hat{t}_{e2e}(thre) = r(thre) \cdot t_{edge} + (1 - r(thre)) \cdot (t_{trans} + t_{cloud}) \quad (7)$$

As EdgeFM does not need users to provide data annotations, we adopt the predictions of FM as ground truth to compute the accuracy as the estimated accuracy  $acc(thre)$ .

At runtime, EdgeFM performs analysis on the estimated accuracy-latency space based on the threshold-searching table and the priority of user demands. For example, if the priority of accuracy is higher than the inference latency, EdgeFM will select the smallest  $thre$  to satisfy not exceeding the constraint of accuracy degradation. If the inference latency has a higher priority, EdgeFM will select the largest  $thre$  to ensure the estimated end-to-end latency is lower than the latency constraint as follows:

$$\max_{thre \in (0,1)} thre \quad \text{s.t.} \quad \hat{t}_{e2e}(thre) \leq L_{app} \quad (8)$$

where  $L_{app}$  is the end-to-end latency constraint, which is specified by the applications. At runtime, network transmission time  $t_{trans}$  can be efficiently updated by estimating the real-time network bandwidth:  $t_{trans} = \frac{Dim}{B(t)}$ , where  $Dim$  is the dimension of samples,  $B(t)$  is the real-time estimated network bandwidth. The estimation of the network has been extensively studied in prior research [31]. Our approach is compatible with the most widely used techniques in this field. Based on Equation 7 and Equation 8, we can obtain the optimal edge-side processing proportion  $r(thre)$  at the current network bandwidth. The relationship between  $r(thre)$  and  $thre$  can be queried from the threshold-searching table with negligible latency. Therefore, EdgeFM can adapt to the network variation through dynamic adjusting its threshold at runtime.

## 5.4 System Implementation

**5.4.1 Edge-cloud Implementation.** We implement EdgeFM on a desktop server (Intel i9-12900K CPU with two NVIDIA RTX 3090 GPU) and two NVIDIA edge platforms, Jetson Nano and Jetson AGX Xavier. The network connection and data transmission parts in EdgeFM are developed via TCP socket API. We use Traffic Control in Linux to simulate different network conditions and iPerf tool [28] to measure the network bandwidth at regular one-second intervals.

**5.4.2 Foundation Models.** We implement EdgeFM on two FMs, ImageBind [18] and CLIP [52]. We adopt the image and audio modalities of ImageBind for evaluation in this work. For CLIP, we use the CLIP-L/14 version, which reports the highest performance among CLIP series. For vision-based tasks, we use CLIP and the vision

branch of ImageBind as FMs for evaluation. As CLIP only supports vision modality, we only use ImageBind for the audio recognition task.

**5.4.3 Prompt Setting.** Both ImageBind and CLIP require a prompt to convert the single class name in a natural language manner into a textual description. The prompt for HAR task is set to “a photo of a person doing *CLS*.”, which are the same as CLIP’s setting [52]. For indoor scene recognition and flower recognition, the prompt is “a photo of a *CLS*.”. For audio recognition, we extract the text embeddings from the class name, which is the same as ImageBind’s setting [18].

**5.4.4 Baseline Approaches.** We compare EdgeFM with two types of baselines, including the efficient on-device inference baselines and open-set recognition baselines.

**Efficient On-device Inference Baselines.** We implement several representative on-device NN efficient inference approaches on ImageBind and CLIP for a fair evaluation.

*PersEPhonEE* [35], which is an edge-only NN acceleration approach based on early exit. we implement PersEPhonEE on the two FMs. There are two ways to implement the early-exit classifier for ImageBind and CLIP, including a fully-connected classifier and cosine distance classifier [10], where the latter one is adopted as it performs better in experiments.

*SPINN* [34], which is an edge-cloud collaboration approach integrating model splitting and early-exit techniques. Similarly, we re-implement SPINN on ImageBind and CLIP. We also adopt cosine distance classifier as the early-exit head.

*Cloud-centric*, which is the most widely adopted solution for FM inference [1, 38, 53]. For the cloud-centric approach, we deploy ImageBind and CLIP on the server and offload all the samples to the server for inference.

**Open-set Recognition Baselines.** We also compare the open-set recognition accuracy of EdgeFM with other open-set recognition baselines.

*Semantic-based Approaches.* DUS-VAE [57], ER-ZSAR [8] and VGGishZSL [67] are three typical semantic-based baselines, which are specifically designed for HAR, audio recognition, and image recognition. They all connect the sensor data’s embedding with the semantic embeddings of classes or sentence descriptions generated from language models such as Word2Vec [45] and BERT [14].

*GAN-based approaches.* TF-VAEGAN [46] is a GAN-based open-set recognition approach, which uses a semantic decoder to synthesize features for unseen classes. As current GAN-based approaches mainly focus on vision tasks, they are not used for the comparison on audio-related tasks.

## 6 EVALUATION

### 6.1 Applications and Datasets

We evaluate EdgeFM on three application scenarios, i.e., HAR, robotic semantic sensing, and audio recognition. Table 2 shows the details of the five datasets and their corresponding FM.

#### 6.1.1 Human Activity Recognition.

**Self-collected HAR dataset (SC15).** We collect a real-world HAR dataset in an indoor home setting as shown in Figure 9a. The dataset



**Table 2: Details of datasets, and FMs used in evaluation. CLS is the number of classes.**

Dataset	Tasks	CLS	FMs
SC15	Activity recognition	15	ImageBind/CLIP
UCF101	Activity recognition	101	ImageBind/CLIP
SC40	Indoor scene recognition	40	ImageBind/CLIP
FLO102	Flower recognition	102	ImageBind/CLIP
ESC50	Audio recognition	50	ImageBind

contains 30 subjects and each subject performs 15 activities, including sleeping, sitting on a chair, standing, squatting, playing mobile phone, rummaging cabinet, eating, playing chess, drinking, playing computer, reading books, watering flowers, handwashing, sweeping the floor, and brushing teeth. The total duration of the dataset is about 20 hours, where the sampling rate is 20Hz. The collected video data is split into 2-second recordings, and the middle frame is selected as input samples of EdgeFM, which is the same setting in [18, 52].

**UCF101**[56]. This is a public video-based human activity recognition dataset. We extract the middle frame of each activity video as input samples, which is the same processing way as in [18, 52]. It contains 11,331 images from 101 different categories of human activities. We use the same train-test split as in [52] for a fair evaluation.

### 6.1.2 Robotic Semantic Sensing.

**Self-collected indoor scene dataset (SC40).** We collect a real-world indoor scene dataset, where objects in the environment are allowed to be added dynamically. As shown in Figure 9b, a robot equipped with a camera and edge platform moves randomly in the environment and takes RGB images continuously. We collect 18,295 RGB images in 40 classes in total. During the data collection, we place half of the classes first and place the other later. Current indoor scene datasets like SUN RGB-D [55] do not consider the scenario where users dynamically add objects into the environment, which is the main purpose of our self-collected dataset.

**FLO102** [47]. This is a public dataset that contains 8,189 images from 102 different categories of flowers. We use the same train-test split as in [52] for a fair evaluation.



### 6.1.3 Audio Recognition.

**ESC50** [51]. This is a public audio dataset that contains 2,000 audio clips from 50 different classes of environmental sound. Each audio clip is a 5-sec recording sampled at 44.1 kHz. We use the same train-test split as in [67] for a fair evaluation.

## 6.2 An End-to-End Application

We conduct an end-to-end test by deploying EdgeFM with CLIP as the FM on a mobility robot for semantic sensing. Figure 10a shows the moving trajectory of the robot in a room.

**6.2.1 Adaptability to Network Variation.** We evaluate the adaptability of EdgeFM to network variation. When the robot moves, the network bandwidth fluctuates with time and location, where the lowest and the highest bandwidth are 2 Mbps and 123 Mbps, respectively. We prioritize the execution performance by setting

**(a) HAR testbed setup.****(b) Robot semantic sensing testbed setup.****Figure 9: Two real-world testbed setups.**

the latency bound to 30ms, which can meet real-time requirements for most of the applications [26]. EdgeFM tunes the threshold of model switching according to the bandwidth in real-time, ranging from 0 to 1 with an interval of 0.05.


Figure 10b shows that EdgeFM sets the threshold to a relatively high value ( $\sim 0.99$ ) to ensure that most of the samples are offloaded to the cloud for inference, at high bandwidth conditions (e.g.,  $t \in [50, 200]$ ). When the bandwidth is low (e.g.,  $t \in [20, 50]$ ), EdgeFM sets the threshold to a relatively low value ( $\sim 0.15$ ) to make most data processed on the edge while only few samples are offloaded to the cloud. Overall, results show that EdgeFM can successfully adapt to the dynamic network variation by tuning the threshold of model switching at runtime.

**6.2.2 Adaptability to Environment Change.** We further evaluate the adaptability of EdgeFM to the environment change, i.e. both data distribution and interested classes change. To simulate the scenario where items in users' homes change over time, we first add half of the classes into the environment and then add the remaining classes later. We run EdgeFM continuously in an unsupervised manner to evaluate the adaptability to the environment change.


Figure 11 shows the proportion of edge-cloud processed data, the overall accuracy of EdgeFM and the original FM, and the moment when environment change occurs. The grey dashed line increases after the environment change as the original FM has higher accuracy for the second half of classes. The result shows that EdgeFM can adjust the proportion of edge-cloud processed data at runtime to adapt to the environment change. To maintain the overall accuracy close to the original FM's accuracy, the edge processing proportion of EdgeFM decreases from 84.4% to 40.2% after environment change (i.e., the green bar in Figure 11). Compared with traditional close-set approaches, EdgeFM can reduce the efforts of manual labeling significantly.

## 6.3 Overall Performance of EdgeFM

**6.3.1 Comparison with Baselines of Efficient On-device Inference.** We evaluate both the inference latency and accuracy of EdgeFM and



(a) Moving trajectories of the robot in the room.



(b) Variation of network bandwidth and the threshold and inference latency of EdgeFM.

Figure 10: EdgeFM’s setup and system indicators of the end-to-end evaluation.




Figure 11: EdgeFM’s performance when the environment changes in the end-to-end evaluation. X-axis shows the amount of collected data. D1 and D2 denote the two environments.

on-device NN efficient inference baselines. We implement PersEPhonEE [35] and SPINN [34] on two FMs, i.e., ImageBind and CLIP. We set the bandwidth to 55 Mbps for all tests in this evaluation.

The results in Figure 12 and Table 3 show that EdgeFM can achieve up to  $1.52x \sim 2.63x$  end-to-end latency reduction compared with the best baseline approaches for ImageBind, and can also achieve up to  $1.27x \sim 3.22x$  end-to-end latency reduction compared with the best baseline approaches for CLIP. As shown in Figure 12, among these approaches, only the latency of EdgeFM is lower than 60ms under the 55 Mbps bandwidth, which can meet the real-time requirements for most applications [43]. Meanwhile, EdgeFM can

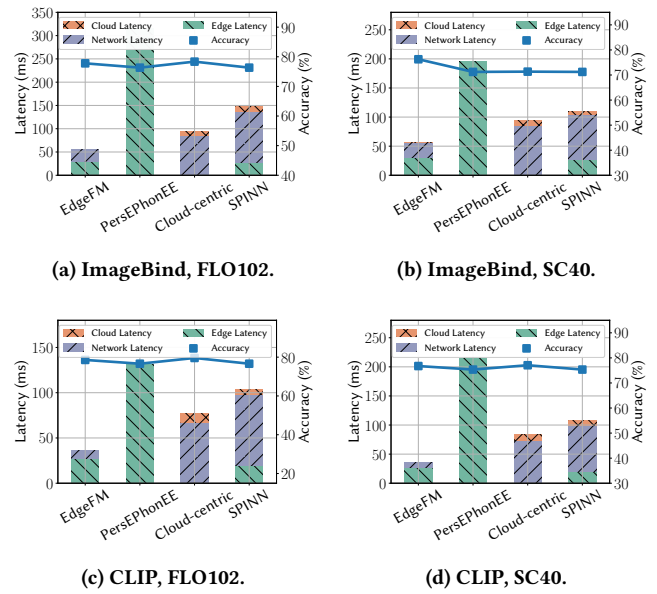


Figure 12: Accuracy and inference latency of EdgeFM and other on-device NN efficient inference systems.

Table 3: The reduced end-to-end latency compared with the best baselines on two edge platforms. N.A. means not applicable since CLIP supports vision data only.

	FM	Device	FLO102	UCF101	SC40	SC15	ESC50
ImageBind	Xavier		1.67x	2.63x	1.65x	1.52x	1.70x
	Nano		1.80x	2.36x	1.84x	1.73x	1.32x
CLIP	Xavier		2.10x	2.20x	2.36x	2.01x	N.A.
	Nano		2.52x	1.27x	3.22x	2.37x	N.A.

achieve higher accuracy than both SPINN and PersEPhonEE on the two FMs. This is because PersEPhonEE adopts an early-exit mechanism to reduce redundant computation. However, early-exit heads on FMs are heavyweight due to the high-dimensional embedding of FMs and deep layers. Although SPINN can offload computation to the cloud, it needs to transmit a large size intermediate embeddings. The size of intermediate embeddings of ImageBind is  $257 \times 1 \times 1280$ , which is much larger than the size of the raw image (i.e.,  $3 \times 224 \times 224$ ).

Another observation is that EdgeFM’s recognition accuracy can even outperform the cloud-centric approach on some certain datasets (see Figure 12b). This is because a dedicated SM trained with abundant data (e.g., more than 800) can outperform FMs, which aligns with the findings shown in Figure 3. However, for more challenging datasets with more diverse data, such as UCF101, the performance of customized small models remains inferior to that of the FMs. This discrepancy also proves the remarkable generalization ability of FMs.

6.3.2 Impact of Network Bandwidth. We evaluate the impact of network bandwidth on EdgeFM. We conduct evaluations under low

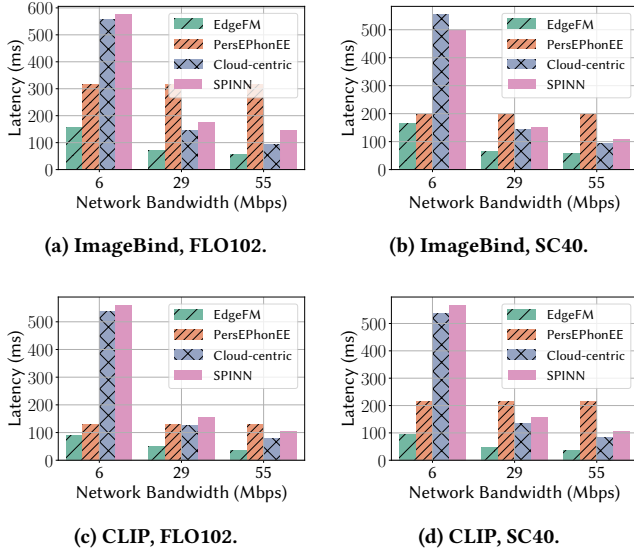


Figure 13: Impact of network bandwidth on EdgeFM.

(6 Mbps), middle (29 Mbps), and high (55 Mbps) network bandwidth, respectively. Figure 13 shows that EdgeFM can achieve up to  $3.5\times$  and  $3.7\times$  end-to-end inference speedup on ImageBind compared with cloud-centric and SPINN under low network conditions (6 Mbps). Under high network bandwidth conditions, this gap narrows, where EdgeFM is still able to achieve up to  $1.7\times$  and  $2.4\times$  end-to-end inference speedup compared with cloud-centric and SPINN under 55 Mbps bandwidth. Results in Figure 13 show that EdgeFM performs better than the existing solutions, especially under low bandwidth conditions.

**6.3.3 Comparison with Open-set Recognition Approaches.** We then evaluate the open-set recognition accuracy and inference latency of EdgeFM and open-set recognition approaches on two edge platforms. The network bandwidth is also set to 55 Mbps for EdgeFM.

As shown in Table 4, EdgeFM achieves  $26.7\%$  higher accuracy on average than GAN-based approaches. Compared with semantic-based approaches, EdgeFM achieves  $21.2\%$  and  $21.7\%$  higher accuracy than DUS-VAE and ES-ZSAR, respectively. Since EdgeFM adopts lightweight small models on the edge, the end-to-end inference latency of EdgeFM is lower than the baselines. For Jetson Nano, EdgeFM reduces the latency by  $1.73\times$  and  $1.98\times$  compared with the best baseline on FLO102 and UCF101 datasets respectively.

We also implement EdgeFM on the audio branch of ImageBind and compare it with baselines on the audio recognition task. EdgeFM achieves  $34.3\%$  accuracy gain and  $3.22\times$  inference latency reduction on Jetson Nano compared with VGGishZSL on ESC dataset. VGGishZSL adopts VGG19 [54] as the audio feature extractor, where its FLOPS and parameters are  $10\times$  larger than the lightweight, small model used in EdgeFM. However, the inference latency of EdgeFM is slower than VGGishZSL  $12.9\text{ms}$  on Xavier. This is because the strong computing power of Jetson AGX Xavier’s GPU makes the parallel computation highly efficient, which causes the latency to be not equivalent to the FLOPS and parameters.

Table 4: Overall accuracy(%) and latency of EdgeFM compared with other open-set recognition approaches.

Dataset	Approach	Acc	Xavier	Nano
FLO102	TF-VAEGAN [46]	62.5%	54.2 ms	108.1 ms
	DUS-VAE [57]	62.1%	57.7 ms	110.1 ms
	<b>EdgeFM</b>	<b>83.3%</b>	<b>44.6 ms</b>	<b>62.2 ms</b>
UCF101	TF-VAEGAN [46]	41.0%	53.9 ms	107.1 ms
	ER-ZSAR [8]	51.8%	87.9 ms	424.8 ms
	<b>EdgeFM</b>	<b>73.5%</b>	<b>42.7 ms</b>	<b>54.0 ms</b>
ESC50	VGGishZSL [67]	33.0%	42.2 ms	217.2 ms
	<b>EdgeFM</b>	<b>67.3%</b>	<b>55.1 ms</b>	<b>67.5ms</b>

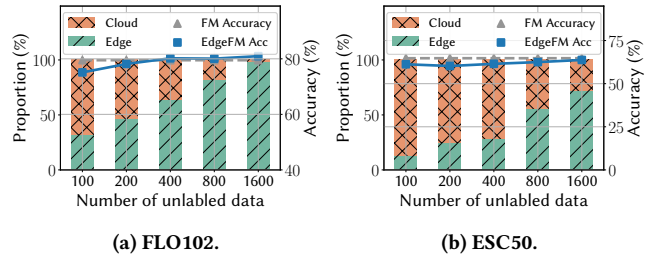


Figure 14: Proportion of the data processed by the customized small model on edge and FM on the cloud.

## 6.4 Understanding EdgeFM’s Performance

**6.4.1 Proportion of the Data Processed on Edge.** We assess the proportion of the data processed by the customized small model on edge and the FM on cloud in EdgeFM, where we use CLIP as the FM in the experiments. Figure 14 shows that the proportion of the edge processed data increases when more data is collected. The proportion of data processed on edge increases from  $31.1\%$  to  $63.5\%$  when the collected data increases from 100 to 400. The proportion can increase up to  $97.3\%$  when collecting 1600 samples in the environment, which means only  $2.7\%$  data are required to be uploaded to the cloud for inference, thus can reduce end-to-end latency compared with cloud-centric solutions. Moreover, our dynamic model switching strategy can keep the overall accuracy always close to the original FM.

**6.4.2 Effectiveness of Semantic-driven Customization.** We compare our semantic-driven customization with the vanilla KD and fine-tuning with the hard pseudo label (FT). The vanilla KD [21] adopts the standard KL divergence to minimize the embedding gap between FMs and small models without using the pseudo text embeddings from FMs. FT refers to adopting the hard pseudo label predicted by FMs as ground truth and cross-entropy for distillation. Figure 15a shows that our semantic-driven customization (marked as SDC in the figure) is able to achieve up to  $9.2\%$ ,  $6.1\%$ ,  $4.7\%$ , and  $6.7\%$  accuracy gain for small model performance under the different number of training data. Since hard pseudo labels fail to preserve semantic relationships between categories [27, 61], FT performs inferior to our approach. Vanilla KD fails to leverage






Figure 15: The performance of EdgeFM’s semantic-driven customization.

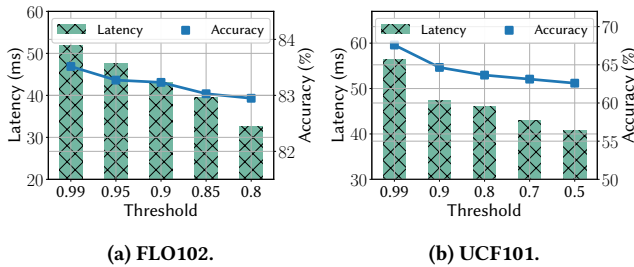


Figure 16: Accuracy-latency trade-off caused by the confidence threshold in EdgeFM.

the knowledge within text embeddings, thus resulting in inferior performance compared to our approach. Figure 15b shows the edge-cloud performance between the three approaches, where we keep the data uploading proportion the same (50% in our setting) for a fair comparison. The result show that EdgeFM’s semantic-driven customization can achieve up to 4.6%, 3.0%, 3.0%, 2.4%, and 3.4% accuracy gain under the different number of training data.

**6.4.3 Trade-off between Accuracy and Latency.** Figure 16 shows the accuracy-latency trade-off caused by the threshold of model switching in EdgeFM. Setting a higher threshold leads to more sensor data offloaded to the cloud, and higher overall accuracy but longer end-to-end inference latency. On the other hand, a lower threshold makes more sensor data processed by the customized small models on the edge and save the network transmission latency, achieving faster end-to-end inference but lower accuracy. Therefore, there is a trade-off between accuracy and latency caused by the model switching thresholds. EdgeFM adjusts the threshold at runtime considering the variation of network bandwidth. The evaluation results can be found in Section 6.2.1.

## 7 DISCUSSION

**Scalability to Other FMs.** This work targets multi-modal FMs based on embedding matching paradigms. EdgeFM deploys CLIP and ImageBind, which are the most popular FMs in the category. There are recent works proposed based on the similar matching method, such as DetCLIPv2 [72] for detection, SAM [32] for segmentation. We envision that EdgeFM’s collaboration between the FMs and specialized small models and iteratively querying specific knowledge from FMs can be extended to other FMs.

**Change of Distribution.** The machine learning models can be formulated as  $\mathcal{Y} = f(\mathcal{X})$ . There is a set of existing works [2, 19, 41] study the change of data distribution  $\mathcal{X}$ . This paper focuses on open-set learning on the edge, i.e., the dynamic change of class set  $\mathcal{Y}$ . In fact, the change of interested class set  $\mathcal{Y}$  can be regarded as a generalized distribution change, including the change of both  $\mathcal{X}$  and  $\mathcal{Y}$ . Our evaluation in Section 6.2.2 implicitly shows that EdgeFM can also well adapt to the distribution changes.

**The Optimal Choice of Small Models.** EdgeFM pre-stores a model pool containing many common small model architectures with diverse accuracy, FLOPS, memory, and latency. This pool also includes architectures through Neural Architecture Search (NAS) [17], such as EfficientNet [59]. Recent studies [40] have investigated the utilization of NAS to look for the best student architecture for the given teacher model during the KD process. They can be integrated with EdgeFM to search for the most suitable small model architecture for a given FM.

**Applications with Labeled Calibration Data.** EdgeFM focuses on the applications without labeled data for calibration to eliminate the overhead of manual labelling. Recent studies have shown that the knowledge in FMs can be better evoked via parameter-efficient fine-tuning (PEFT) approaches [24]. EdgeFM also supports working in the scenario when labeled data is available. In such a case, EdgeFM first uses the labeled data to fine-tune FMs by PEFT on the cloud. Then, the fine-tuned FM can further provide a knowledge base service for small models to query.

**Scalability to Other Sensor Modalities.** EdgeFM supports other time-series sensor data such as video, audio, and IMU. FMs used by EdgeFM, i.e., ImageBind, can support diverse sensor modalities with the corresponding pre-trained encoders. The techniques for video streaming such as frame filtering [9], can also be integrated with EdgeFM to improve the efficiency of processing time-series vision data.

## 8 CONCLUSION

This paper proposes EdgeFM, a novel edge-cloud cooperative system which empowers embedded systems with open-set recognition ability by leveraging FMs for selective knowledge query and edge model customization. EdgeFM maintains the overall performance always close to the FM by dynamic model switching. Extensive experiments show that EdgeFM can reduce the end-to-end latency up to 3.2x and achieve 34.3% accuracy increase compared with the baseline.

## 9 ACKNOWLEDGEMENT

This paper is supported in part by the Research Grants Council (RGC) of Hong Kong under Collaborative Research Fund (CRF) grants C4072-21G and C4034-21G, General Research Fund (GRF) 14214022, Faculty of Engineering of The Chinese University of Hong Kong under Direct Grant 4055167, National Science Foundation of China (NSFC) under Young Scientists Fund 62202407.

## REFERENCES

- [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. 2022. Do As I Can and Not As I Say: Grounding Language in Robotic Affordances. In *arXiv preprint arXiv:2204.01691*.

- [2] Ali Akbari and Roozbeh Jafari. 2019. Transferring activity recognition models for new wearable sensors with deep generative domain adaptation. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*. 85–96.
- [3] Romil Bhardwaj, Zhengxu Xia, Ganesh Ananthanarayanan, Junchen Jiang, Yuan-chao Shu, Nikolaos Karianakis, Kevin Hsieh, Paramvir Bahl, and Ion Stoica. 2022. Eky: Continuous learning of video analytics models on edge compute servers. In *19th USENIX Symposium on Networked Systems Design and Implementation*. 119–135.
- [4] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- [6] Qingqing Cao, Prerna Khanna, Nicholas D Lane, and Aruna Balasubramanian. 2022. MobiVQA: Efficient On-Device Visual Question Answering. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 2 (2022), 1–23.
- [7] Lingjiao Chen, Matei Zaharia, and James Zou. 2023. FrugalGPT: How to Use Large Language Models While Reducing Cost and Improving Performance. *arXiv preprint arXiv:2305.05176* (2023).
- [8] Shizhe Chen and Dong Huang. 2021. Elaborative rehearsal for zero-shot action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13638–13647.
- [9] Tiffany Yu-Han Chen, Lenin Ravindranath, Shuo Deng, Paramvir Bahl, and Hari Balakrishnan. 2015. Glimpse: Continuous, real-time object recognition on mobile devices. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. 155–168.
- [10] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. 2019. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232* (2019).
- [11] Google Cloud and Vertex AI. 2021. [Online]. Available: <https://cloud.google.com/vertex-ai>.
- [12] Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. 2022. Enabling Multimodal Generation on CLIP via Vision-Language Knowledge Distillation. In *Findings of the Association for Computational Linguistics: ACL 2022*. Association for Computational Linguistics, Dublin, Ireland, 2383–2395.
- [13] Deloitte. 2019. Edge AI chip shipments by device worldwide 2020 and 2024. <https://www.statista.com/statistics/1084670/edge-ai-chips-shipment-worldwide/>.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186.
- [15] Shuya Ding, Zhe Chen, Tianyue Zheng, and Jun Luo. 2020. RF-net: A unified meta-learning framework for RF-enabled one-shot human activity recognition. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 517–530.
- [16] Utsav Drolia, Katherine Guo, Jiaqi Tan, Rajeev Gandhi, and Priya Narasimhan. 2017. Cachier: Edge-caching for recognition applications. In *2017 IEEE 37th International Conference on Distributed Computing Systems*. IEEE, 276–286.
- [17] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2019. Neural architecture search: A survey. *The Journal of Machine Learning Research* 20, 1 (2019), 1997–2017.
- [18] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15180–15190.
- [19] Taesik Gong, Yeonsu Kim, Jinwoo Shin, and Sung-Ju Lee. 2019. Metasense: few-shot adaptation to untrained conditions in deep mobile sensing. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 110–123.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*.
- [22] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. 2019. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1314–1324.
- [23] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [24] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Sheng Wang, Lu Wang, Weizhu Chen, et al. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- [25] Zhiming Hu, Ning Ye, and Iqbal Mohamed. 2022. mmSampler: Efficient Frame Sampler for Multimodal Video Retrieval. *Proceedings of Machine Learning and Systems* 4 (2022), 153–171.
- [26] Kai Huang and Wei Gao. 2022. Real-time neural network inference on extremely weak devices: agile offloading with explainable AI. In *Proceedings of the 28th Annual International Conference on Mobile Computing and Networking*. 200–213.
- [27] Md Tamzeed Islam and Shahriar Nirjon. 2019. Soundsemantics: exploiting semantic knowledge in text for embedded acoustic event classification. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*. 217–228.
- [28] Dugan Jon, Elliott Seth, Bruce A. Mah, Poskanzer Jeff, and Prabhu Kaustubh. 2014. iPerf. <https://software.es.net/iperf/>.
- [29] Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. 2017. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. *ACM SIGARCH Computer Architecture News* 45, 1 (2017), 615–629.
- [30] Mehrdad Khani, Ganesh Ananthanarayanan, Kevin Hsieh, Junchen Jiang, Ravi Netravali, Yuanchao Shu, Mohammad Alizadeh, and Victor Bahl. 2023. RECL: Responsive Resource-Efficient Continuous Learning for Video Analytics. In *20th USENIX Symposium on Networked Systems Design and Implementation*. 917–932.
- [31] Minkyong Kim and Brian Noble. 2001. Mobile network estimation. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*. 298–309.
- [32] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. Segment Anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4015–4026.
- [33] Ikki Kishida, Hong Chen, Masaki Baba, Jiren Jin, Ayako Amma, and Hideki Nakayama. 2021. Object recognition with continual open set domain adaptation for home robot. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1517–1526.
- [34] Stefanos Laskaridis, Stylianos I Venieris, Mario Almeida, Ilias Leontiadis, and Nicholas D Lane. 2020. SPINN: synergistic progressive inference of neural networks over device and cloud. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–15.
- [35] Ilias Leontiadis, Stefanos Laskaridis, Stylianos I Venieris, and Nicholas D Lane. 2021. It's always personal: Using early exits for efficient on-device CNN personalisation. In *Proceedings of the 22nd International Workshop on Mobile Computing Systems and Applications*. 15–21.
- [36] Min Li, Yu Li, Ye Tian, Li Jiang, and Qiang Xu. 2021. AppealNet: An efficient and highly-accurate edge/cloud collaborative architecture for DNN inference. In *2021 58th ACM/IEEE Design Automation Conference*. 409–414.
- [37] Xiang Li, Xin Jiang, Xuying Meng, Aixin Sun, and Yequan Wang. 2023. FreeLM: Fine-Tuning-Free Language Model. *arXiv preprint arXiv:2305.01616* (2023).
- [38] Xiwen Liang, Yangxin Wu, Jianhua Han, Hang Xu, Chunjing Xu, and Xiaodan Liang. 2022. Effective adaptation in multi-task co-training for unified autonomous driving. *Advances in Neural Information Processing Systems* 35 (2022), 19645–19658.
- [39] Edo Liberty, Zohar Karmin, Bing Xiang, Laurence Rouesnel, Baris Coskun, Ramesh Nallapati, Julio Delgado, Amir Sadoughi, Yury Astashonok, Piali Das, et al. 2020. Elastic machine learning algorithms in amazon sagemaker. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 731–737.
- [40] Yu Liu, Xuhui Jia, Mingxing Tan, Raviteja Vemulapalli, Yukun Zhu, Bradley Green, and Xiaogang Wang. 2020. Search to distill: Pearls are everywhere but not the eyes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7539–7548.
- [41] Wenjie Luo, Zhenyu Yan, Qun Song, and Rui Tan. 2021. Phyaug: Physics-directed data augmentation for deep sensing model transfer in cyber-physical systems. In *Proceedings of the 20th International Conference on Information Processing in Sensor Networks*. 31–46.
- [42] Akhil Mathur, Anton Isopoussu, Nadia Berthouze, Nicholas D Lane, and Fahim Kawar. 2019. Unsupervised domain adaptation for robust sensory systems. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*. 505–509.
- [43] Zili Meng, Tingfeng Wang, Yixin Shen, Bo Wang, Mingwei Xu, Rui Han, Honghao Liu, Venkat Arun, Hongxin Hu, and Xue Wei. 2023. Enabling High Quality Real-Time Communications with Adaptive Frame-Rate. In *20th USENIX Symposium on Networked Systems Design and Implementation*. 1429–1450.
- [44] Benjamin J Meyer and Tom Drummond. 2019. The importance of metric learning for robotic vision: Open set recognition and active learning. In *2019 International Conference on Robotics and Automation*. IEEE, 2924–2931.
- [45] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality.



- Advances in Neural Information Processing Systems* 26 (2013).
- [46] Sanath Narayan, Akshita Gupta, Fahad Shahbaz Khan, Cees GM Snoek, and Ling Shao. 2020. Latent embedding feedback and discriminative features for zero-shot classification. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII* 16. Springer, 479–495.
- [47] Maria-Elena Nilsback and Andrew Zisserman. 2008. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*. IEEE, 722–729.
- [48] Wei Niu, Xiaolong Ma, Sheng Lin, Shihao Wang, Xuehai Qian, Xue Lin, Yanzhi Wang, and Bin Ren. 2020. Patdnn: Achieving real-time dnn execution on mobile devices with pattern-based weight pruning. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 907–922.
- [49] Xiaomin Ouyang, Xian Shuai, Jiayu Zhou, Ivy Wang Shi, Zhiyuan Xie, Guoliang Xing, and Jianwei Huang. 2022. Cosmo: contrastive fusion learning with small data for multimodal human activity recognition. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*. 324–337.
- [50] Eunhyeok Park, Dongyoung Kim, Soobeom Kim, Yong-Deok Kim, Gunhee Kim, Sungroh Yoon, and Sungjoo Yoo. 2015. Big/little deep neural network for ultra low power inference. In *2015 International Conference on Hardware/Software Codesign and System Synthesis*. IEEE, 124–132.
- [51] Karol J Piczak. 2015. ESC: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*. 1015–1018.
- [52] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [53] Dhruv Shah, Błażej Osiński, Sergey Levine, et al. 2023. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on Robot Learning*. PMLR, 492–504.
- [54] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [55] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. 2015. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 567–576.
- [56] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012).
- [57] Hongzu Su, Jingjing Li, Zhi Chen, Lei Zhu, and Ke Lu. 2022. Distinguishing unseen from seen for generalized zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7885–7894.
- [58] Ximeng Sun, Pengchuan Zhang, Peizhao Zhang, Hardik Shah, Kate Saenko, and Xide Xia. 2023. DIME-FM: Distilling Multimodal and Efficient Foundation Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15521–15533.
- [59] Mingxing Tan and Quoc Le. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*. PMLR, 6105–6114.
- [60] MLC team. 2023. *MLC-LLM*. <https://github.com/mlc-ai/mlc-llm>
- [61] Catherine Tong, Jinchen Ge, and Nicholas D Lane. 2021. Zero-shot learning for imu-based activity recognition using video embeddings. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (2021), 1–23.
- [62] Sathish Vallachira, Michal Orkisz, Mikael Norrlöf, and Sachit Butail. 2019. Data-driven gearbox failure detection in industrial robots. *IEEE Transactions on Industrial Informatics* 16, 1 (2019), 193–201.
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).
- [64] Yiding Wang, Kai Chen, Haisheng Tan, and Kun Guo. 2023. Tabi: An Efficient Multi-Level Inference System for Large Language Models. In *Proceedings of the Eighteenth European Conference on Computer Systems*. 233–248.
- [65] Yixuan Wei, Han Hu, Zhenda Xie, Zheng Zhang, Yue Cao, Jianmin Bao, Dong Chen, and Baining Guo. 2022. Contrastive learning rivals masked image modeling in fine-tuning via feature distillation. *arXiv preprint arXiv:2205.14141* (2022).
- [66] Jiayang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. 2016. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4820–4828.
- [67] Huang Xie and Tuomas Virtanen. 2021. Zero-shot audio classification via semantic embeddings. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021), 1233–1242.
- [68] Huatao Xu, Pengfei Zhou, Rui Tan, Mo Li, and Guobin Shen. 2021. Limu-bert: Unleashing the potential of unlabeled data for imu sensing applications. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*. 220–233.
- [69] Mengwei Xu, Mengze Zhu, Yunxin Liu, Felix Xiaozhu Lin, and Xuanzhe Liu. 2018. Deepcache: Principled cache for mobile deep vision. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 129–144.
- [70] Bufang Yang, Wenxuan Wu, Yitian Liu, and Hongxing Liu. 2022. A novel sleep stage contextual refinement algorithm leveraging conditional random fields. *IEEE Transactions on Instrumentation and Measurement* 71 (2022), 1–13.
- [71] Bufang Yang, Xilin Zhu, Yitian Liu, and Hongxing Liu. 2021. A single-channel EEG based automatic sleep stage classification method leveraging deep one-dimensional convolutional neural network and hidden Markov model. *Biomedical Signal Processing and Control* 68 (2021), 102581.
- [72] Lewei Yao, Jianhua Han, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, and Hang Xu. 2023. Detclipv2: Scalable open-vocabulary object detection pre-training via word-region alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 23497–23506.
- [73] Shuochao Yao, Jinyang Li, Dongxin Liu, Tianshi Wang, Shengzhong Liu, Huajie Shao, and Tarek Abdelzaher. 2020. Deep compressive offloading: Speeding up neural network inference by trading edge computation for network latency. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 476–488.
- [74] Mu Yuan, Lan Zhang, Fengxiang He, Xueting Tong, and Xiang-Yang Li. 2022. Infi: end-to-end learnable input filter for resource-efficient mobile-centric inference. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*. 228–241.
- [75] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung Ho Bae, Seungkyu Lee, and Choong Seon Hong. 2023. Faster Segment Anything: Towards Lightweight SAM for Mobile Applications. *arXiv preprint arXiv:2306.14289* (2023).
- [76] Zhihe Zhao, Kai Wang, Neiwon Ling, and Guoliang Xing. 2021. Edgempl: An automl framework for real-time deep learning on the edge. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*. 133–144.