Open-TransMind: A New Baseline and Benchmark for 1st Foundation Model Challenge of Intelligent Transportation

Yifeng Shi[1][*][†]   Feng Lv[1*]   Xinliang Wang[1*]   Chunlong Xia[1*]
Shaojie Li[1*]   Shujie Yang[1]   Teng Xi[1]   Gang Zhang[1]

[1]Baidu Inc.

Abstract

*With the continuous improvement of computing power and deep learning algorithms in recent years, the foundation model has grown in popularity. Because of its powerful capabilities and excellent performance, this technology is being adopted and applied by an increasing number of industries. In the intelligent transportation industry, artificial intelligence faces the following typical challenges: few shots, poor generalization, and a lack of multi-modal techniques. Foundation model technology can significantly alleviate the aforementioned issues. To address these, we designed the **1st Foundation Model Challenge**, with the goal of increasing the popularity of foundation model technology in traffic scenarios and promoting the rapid development of the intelligent transportation industry. The challenge is divided into two tracks: all-in-one and cross-modal image retrieval. Furthermore, we provide a new baseline and benchmark for the two tracks, called Open-TransMind. According to our knowledge, Open-TransMind is the first open-source transportation foundation model with multi-task and multi-modal capabilities. Simultaneously, Open-TransMind can achieve state-of-the-art performance on detection, classification, and segmentation datasets of traffic scenarios. Our source code is available at* `https://github.com/Traffic-X/Open-TransMind`.

1. Introduction

In recent years, the foundation model has gained a lot of attention in the artificial intelligence community and is one of the key areas for future advancement. Large-scale neural networks that have been trained on a massive amount of data and may be utilized to tackle complex problems are often the building blocks of foundation models. With deep learning's ongoing development, larger and larger deep learn-
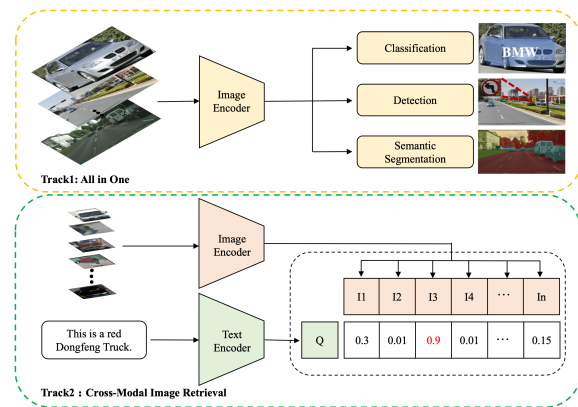

Figure 1. **Top:** Track 1 consists of three tasks including classification, detection and semantic segmentation. **Bottom:** Track 2 is a text-image retrieval task, where Q represents the feature obtained by text encoding, and I represents the images to be queried.

ing models are being trained using an increasing amount of computing resources. This has sped up the development of foundation model technology, which is now finding numerous uses in areas like speech recognition, image and video processing, machine translation, and natural language processing.

Compared to small-scale models, foundation models have many advantages. Foundation models can learn more from larger-scale datasets and perform better overall because of their higher computing power and more parameters. Furthermore, foundation models are capable of learning from data and adapting to new tasks, so they can better adapt to various application contexts. However, foundation models have some drawbacks, such as deployment and inference costs, so many works have conducted extensive research in this area [3, 18, 19].

Intelligent transportation is a field that integrates artificial intelligence technology and the transportation sector, focusing especially on autonomous driving, vehicle-

---

[*]indicates equal contribution

[†]Yifeng Shi(shiyifeng@baidu.com is the corresponding author.)

infrastructure cooperation, traffic management, traffic information services, and other related topics. The following challenges face artificial intelligence in the field of intelligent transportation:

• Scenes with few shot problems have a weak effect;

• The data input, learning capacity, and generalization capacity of small-scale models are constrained;

• The pace of multi-modal technology development is slow, and industry applicability is restricted.

To address these issues, we designed the **1st Foundation Model Challenge** with the goal of increasing the popularity of foundation model technology in traffic scenarios. This activity can help strengthen the application of artificial intelligence technology in the transportation field and promote the rapid development of the intelligent transportation industry. The foundation model challenge consists of two tracks, as shown in Figure 1:

Track 1: All in One. When designing well-crafted neural network structures and loss functions, multi-tasking can significantly improve model generalization. Using only one task's data for training may result in overfitting due to data noise in specific tasks. Unified large models will combine datasets from multiple tasks into unified training, allowing them to average noise from different jobs and thus improve their ability to learn better features. As a result, this track hopes that the unified foundation model will outperform multiple single-task models at the same time, allowing it to explore its upper bound on capabilities.

Track 2: Cross-Modal Image Retrieval. In traffic scenarios, high-performance image retrieval technology is critical for security management. The traditional image retrieval method typically begins with image attribute recognition and then achieves retrieval capability by comparing the image to the expected attribute. However, as multi-modal foundation model technology has advanced, text and image representation unification and modality transformation have become commonplace. The use of this capability can further improve the accuracy and flexibility of image retrieval. As a result, in this track, we created a text-image retrieval dataset that contestants can use to investigate multi-modal technology research, thereby improving the accuracy of text retrieval from images.

In this paper, we provide our baseline methodology, Open-TransMind, and a detailed benchmark introduction, including dataset compositions and evaluation strategies for these two tracks. Open-TransMind is, as far as we are aware, the first open-source foundation model for intelligent transportation and achieve state-of-the-art performance on detection, classification, and segmentation datasets of traffic scenarios. The open-source URL is https://github.com/Traffic-X/Open-TransMind. In the following chapters, we will introduce related work in Chapter 2, benchmark details, including dataset setup and evalua-

tion strategies for each track, in Chapter 3, our baselines in Chapter 4, experiment analysis and results in Chapter 5, and the article conclusion in Chapter 6.

2. Related Work

2.1. Intelligent Transportation Public Datasets

Publicly available datasets have played a critical role in advancing research in the field of intelligent transportation systems and autonomous driving by providing researchers with access to high-quality, diverse data to develop and test algorithms and systems. These datasets are divided into three task types: object detection, image classification, and semantic segmentation.

Object Detection. KITTI [9] includes a large number of real-world scenes captured from moving vehicles, with data from various sensors such as cameras and LiDAR. Annotations of 2D and 3D bounding boxes for various types of objects, such as cars, pedestrians, and cyclists, are included in the dataset. TT100K [25] is a large-scale traffic sign detection benchmark comprised of 100,000 images of road scenes and 30,000 traffic sign instances each annotated with a class label, bounding box, and pixel mask. Other datasets, in addition to frontal view with sensors mounted on the vehicle, focus on roadside perception. Rope3D [20] creates designs for 3D object detection tasks in the roadside view. DAIR-V2X [22] is a 71,254 frame dataset of image and point cloud data designed for comprehensive research on 3D object detection in vehicle-infrastructure cooperative autonomous driving.

Image Classification. The PA100K dataset [13] includes annotations for 26 different pedestrian attributes, such as gender, age group, clothing color, etc. The annotations are provided in the form of binary labels, indicating the presence or absence of each attribute in the corresponding pedestrian image. The Stanford Cars dataset [12] is a large-scale collection of car images collected for fine-grained classification. It includes 16,185 images of 196 different car models taken in a variety of settings and conditions.

Semantic Segmentation. The Cityscapes dataset [4] contains 5,000 high-resolution images captured in various weather and lighting conditions from 50 cities and provides pixel-level annotations for semantic segmentation tasks. Apolloscape [14] includes over 140,000 high-resolution images with 2D pixel-wise and 3D point-wise semantic annotation.

However, annotations for multiple tasks within a single image are not always present in the datasets mentioned above. The BDD100K dataset [21] provides annotations for multiple tasks, including object detection, semantic segmentation, lane detection, and so on, to aid research on multi-task learning in the field of autonomous driving and

Task	Datasets	Perspective			Modality	
		Roadside	Vehicle	Pedestrian	Visual	Language
Classification	Stanford Cars	✓	✓	✓	✓	
Detection	TT100K	✓		✓	✓	
Segmentation	BDD100K	✓			✓	
Text-Image Retrieval	PA100K+BIT-Vehicle+Web data	✓	✓	✓	✓	✓

Table 1. Overview of the used datasets for Track1 and Track2.

intelligent transportation systems.

2.2. Foundation Models

Large, pre-built neural network models that have been trained on massive amounts of data are referred to as pre-trained foundation models. When applied to downstream tasks, foundational models can exhibit strong generalization capabilities with minimal customization. Foundation models are classified into three types based on their modality: LLMs (Large Language Models), LVMs (Large Vision Models), and VLMs (Large Vision-Language Models).

LLMs. Large language models like BERT [6] and GPT-3 [2] have transformed natural language processing (NLP) by producing astounding results on a wide range of NLP tasks like language understanding, question answering, summarization, translation, and so on. LLMs usually adopt unsupervised learning techniques to pre-train transformer-based models [17] on massive amounts of text data.

LVMs. Previous large vision models were frequently based on contrastive learning on CNN-based architecture. However, given the enormous impact of transformer-based models in the field of NLP, researchers wonder whether they can apply the transformer concept to image data. ViT [8] proposes modifying the transformer architecture to process images as a series of patches. Because of its large model capacity and generalizing capability, this transformer-based architecture is widely used in LVMs. MAE [10] applies the concept of reconstruction from BERT to the domain of computer vision. UFO [19] proposes a feature optimization paradigm that pretrains a large model on multiple tasks and then trims the model to a moderate size for deployment on specific tasks using NAS and Task-MoE.

VLMs. VLMs can perform both language and image-related tasks by combining LLMs and LVMs, such as image captioning, text-guided image generation, visual question answering, and so on. CLIP [16] is made up of a visual encoder and a text encoder, and it trains the model using contrastive learning on a large dataset of image-caption pairs. Following research has concentrated on the interaction of the image encoder and the text encoder. Some studies treat the two encoders as separate and independent models, while others [11] try to integrate them into a single, unified model. These studies seek to improve the model's

performance by investigating various methods for combining and utilizing the visual and semantic features learned by the image and text encoders.

3. Benchmark

According to the perspective of data collection, the research content of intelligent transportation mainly includes vehicle perspective, roadside perspective, and pedestrian perspective scene tasks. These tasks include image classification, object detection, segmentation, and text-image retrieval, and the modalities include image, text, point cloud, etc. Considering the richness of tasks and resource occupation, we selected four tasks in three categories: roadside, vehicle, and pedestrian perspectives, which include image classification, object detection, semantic segmentation, and text-image retrieval (shown in Table 1). In view of the above four tasks, we divide them into two tracks. Track 1 is composed of three tasks: image classification, object detection, and semantic segmentation. Track 2 is a cross-modal task of text-image retrieval. The following is a detailed benchmark description of Track 1 and Track 2.

3.1. Track 1: All In One

This track consists of three tasks: image classification, object detection, semantic segmentation. To accommodate as many perspectives and tasks as possible, we select three datasets: Stanford Car [12], TT100K [25] and BDD100K [21], which cover the vehicle, roadside, pedestrian perspective, and the three tasks listed above. Based on the chosen datasets, we created a method for multi-task co-evolution. We'll discuss datasets and evaluation measures in the sections below.

Dataset. BDD100K is a diverse visual-driven scene dataset gathered from vehicle perspectives that includes ten different tasks such as detection, segmentation, and tracking. As our benchmark, we chose the semantic segmentation task dataset with 720×1280 resolution; the dataset contains geographic, environmental, and weather diversity, as well as 7,000 and 1,000 finely annotated images for training and testing, respectively. Tsinghua-Tencent 100K (TT100K) is a large-scale traffic-sign detection and classification benchmark obtained from vehicle and pedestrian perspectives, with 100,000 images containing 30,000 traffic
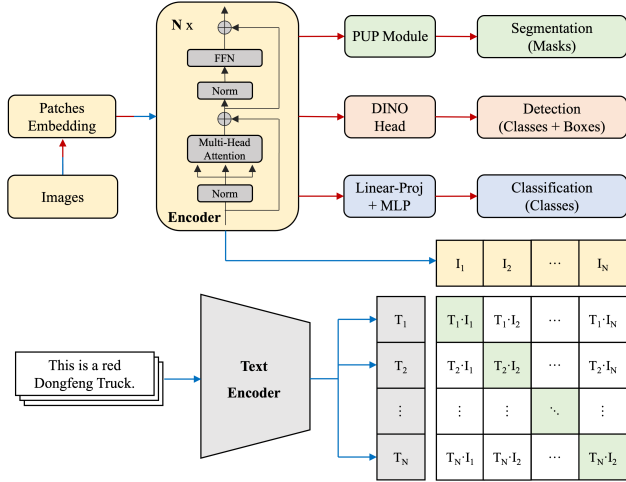
Figure 2. The framework of the large model. The red arrow represents the sub-framework of Track 1, and the blue represents the sub-framework of Track 2.

Figure 3. The structure of Task-MoE with shared FFN and specific FFN.

sign instances annotated with a class label, bounding box, and pixel mask. These images cover a wide range of lighting and weather conditions. We chose traffic-sign detection as our task, which included 6,107 and 3,067 annotated images for training and testing, respectively. Stanford Cars consists of 196 car classes with 360×240 pixels, including 8,144 training images and 8,041 testing images. It is compiled from vehicle, roadside, and pedestrian perspectives. Brand, model, and year are used to categorize items.

Evaluation Metrics. This track hosts a collaborative multi-task optimization competition with classification, detection, and segmentation tasks. As a result, we include not only the evaluation metric for each task, but also a global metric that expresses the overall effect of the competition task. The classification's evaluation metric is Acc. mAP is the detection evaluation metric. The semantic segmentation evaluation metric is mIoU. As the global metric, this track uses the average of all metrics.

3.2. Track 2: Cross-Modal Image Retrieval

This track focuses on multi-modal text-image retrieval task involving two types of transportation participants: pedestrians and vehicles. We generate the images for this track using open-source datasets and internet data, and we used a large language model to generate text for each image. In terms of model design, we use a concise CLIP model and $mAP@k$ as the evaluation metric. Below is a detailed description of the data construction and evaluation metric calculation.

Dataset. For pedestrian data, we use the open-source PA-100K [13], which is the largest dataset for pedestrian attributes recognition task. It contains 90,000 training im-
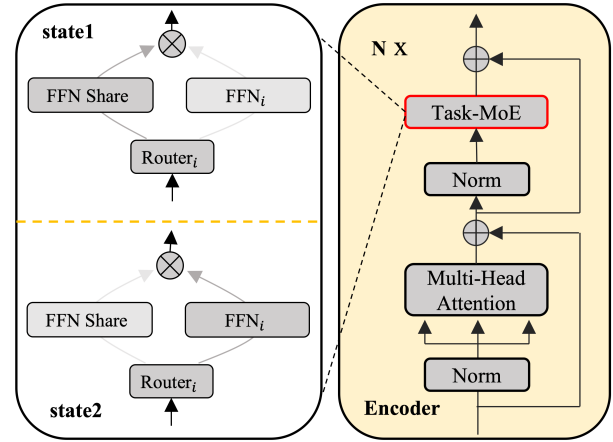
ages and 10,000 testing images with fine-grained attribute annotations such as pedestrian gender, age, upper and lower clothing style, whether or not they wear glasses and so on. Using a large language model, we generate descriptive text based on these attributes. As well as for vehicle data, we use an open-source dataset, BIT-Vehicle [7] to construct the test set. This dataset contains 9,850 vehicle images, with each image annotated with vehicle type attributes. Based on this dataset, we use a proprietary vehicle structured model that has a high accuracy and credibility for vehicle color and brand classification (98.84% for color and 96.48% for brand), to annotate each image with additional vehicle color and brand attributes. After filtering out rare vehicle brand data and manually checking the accuracy of all attributes, we utilize a total of 7,611 images from the dataset, with each image annotated with vehicle color, brand, and type attributes. After statistical analysis, the test set we have constructed contains a total of 11 colors, 65 vehicle brands, and 6 vehicle types. We use these attributes as keywords to crawl 46,117 images from the internet as the training set, and generate corresponding descriptive texts for each image using a large language model based on the attribute category. It is important to note that we have not cleaned the web-scraped training set, which contains noise: certain images may have corresponding text with attribute errors.

After merging pedestrian and vehicle data, Track2 has a total of 137,117 data for training and 17,611 data for testing. Furthermore, in order to facilitate the validation of model capabilities, we randomly select 10,000 data from the training set as the validation set.

Evaluation Metrics. The evaluation metric used in this competition is the mean Average Precision ($mAP@K$), where K is set to 10. mAP is a measure of the accuracy of text-based image retrieval. The calculation of $mAP@K$ is

Method	Backbone	Type	Tasks			Stanford Cars	TT 100K	BDD100K
			cls	det	seg	(acc%)	(mAP%)	(mIoU%)
CAP [1]	xception		✓			95.70	-	-
CABNet [5]	vgg16	single		✓		-	78.00	-
NiseNet [15]	resnet101				✓	-	-	53.52
Open-TransMind	vit-base	single	✓			82.78	-	-
				✓		-	76.27	-
					✓	-	-	52.55
Open-TransMind	vit-base	multi	✓	✓	✓	91.64	76.90	55.13
Open-TransMind	vit-huge	single	✓			95.85	-	-
				✓		-	82.10	-
					✓	-	-	63.41
Open-TransMind	vit-huge	multi	✓	✓	✓	**95.96**	**83.24**	**64.80**

Table 2. Comparing our Open-TransMind approach based on the vit-huge backbone with other state-of-the-art methods and comparing the performance of single-task trained models and multi-task trained models based on the Open-TransMind with vit-base backbone on different task datasets.

as follows:

$$mAP@K = \frac{1}{m} \times \sum_{i=1}^{K} p(i) * \Delta r(i)$$

where m is the total number of text queries in the evaluation set, $p(i)$ represents the precision of the $topi$ retrieved results, and $\Delta r(i)$ is calculated as:

$$\Delta r(i) = r(i) - r(i-1)$$

where $r(i)$ is the recall of the $topi$ retrieved results and $r(0) = 0$.

4. Baseline

In intelligent transportation scenarios, the types of tasks mainly involve multiple fields such as classification, detection, and segmentation. We can usually solve the corresponding problems by building a single-task model, but the generalization ability of the model will be poor due to the small size of the data. Considering the similarity between different tasks in the traffic scene, we propose a multi-task co-evolution model to promote mutual learning between different tasks, of course, there may be conflicts. In addition, traditional image retrieval methods usually use image attribute recognition to achieve retrieval capabilities. However, introducing cross-modal capabilities based on existing models, such as image-text retrieval, can further improve the flexibility of image retrieval applications. Therefore, we can use a multi-task co-evolutionary model to build an image model for text retrieval. Finally, we combine multi-task co-evolution and multi-modal text-image retrieval into a model called Open-TransMind.

4.1. Task-MoE

Our Open-TransMind model is built based on the UFO [19] multi-task and multi-path training framework. Due to the differences in the visual representation of the data used by different tasks, there may be conflicts between tasks. Further, we can alleviate conflicts by introducing Task-MoE. As shown in the Figure 3, the multi-path FFN module is set in the encoding module. Each task has two different path options, that is, to choose shared FFN (state1 in Figure 3) or specific FFN (state2 in Figure 3). All tasks will update the parameters of shared FFN, and specific tasks will only update the specific FFN parameters to achieve collaborative optimization between different tasks, and greatly improve the generalization ability of the model.

4.2. Model Design for All in One

The framework of the large multi-task model for this track is shown in Figure 2. We select the backbone of the transformer architecture as the multi-task feature encoder. Based on this common feature, we respectively perform feature decoding for image classification, object detection, and semantic segmentation tasks and predict the results of corresponding tasks. The feature decoder for classification tasks consists of a linear-projection layer and a full-connection layer. The function of the linear-projection layer is to provide a layer of isolation between classification and multi-task features; the object detection decoder refers to the detection head of the DINO [23] model, and the semantic segmentation decoder uses the progressive upsampling module of the SETR [24] model.

4.3. Model Design for Cross-Modal Image Retrieval

The text-image retrieval model of this track is implemented using the same method as CLIP, and the overall

Training Method	Test	Val
training from scratch	21.25	53.84
CLIP	28.11	16.44
CLIP finetune	50.96	72.26
CLIP finetune + All in One off-the-shelf	**53.54**	**72.56**

Table 3. The quantitative results of the cross-modal text-image retrieval task on the test and validation sets we constructed.

framework is shown in the lower half of Figure 2. We use the same transformer structure as Track 1 as the visual encoder, and use a text encoder to complete feature extraction of text data. During the training process, visual encoders and text encoders calculate visual and text features respectively, and achieve comparative learning of picture text pairs through contrastive training.

5. Experiments

5.1. All In One

Implementation Details. We chose vit-huge and vit-base models of transformers with different sizes as the backbone of Open-TransMind and set different output heads for the three different tasks of detection, segmentation, and classification. We merged three tasks and used a unified hyperparameter, such as the AdamW optimizer with an initial learning rate of 0.0001 and a weight decay of 0.0005. We adopt CosineAnnealingLR to drop the learning rate and implement different data augmentation methods for different tasks, including random resize, random cropping, random scaling in the range of 0.5 to 2.0, and random horizontal flipping, etc. Finally, we trained 120 epochs based on the ImageNet pre-training model with a batch size of 8 on eight A100 GPUs.

Evaluation Results. Based on the above experimental configuration, we conducted multiple sets of comparative ablation experiments. We chose the state-of-the-art methods described on https://paperswithcode.com/ as baselines for comparison in order to confirm the efficacy of our approach, such as CAP, CABNet, and NiseNet, which are currently the best methods for classification, detection, and segmentation datasets of this track. To begin, we compared the performance of multi-task and single-task large models using Open-TransMind based on the vit-base backbone to validate the benefits of multi-task learning. As shown in Table 2, the multi-task joint training large model has significant advantages in the two tasks of segmentation and classification, which are respectively positive by 2.58% and 8.86%, and the detection task is slightly positive. Multi-task joint training of large models has been shown to have higher generalization ability than single-task large models. Second, we evaluated state-of-the-art methodologies against our Open-TransMind method built on the vit-
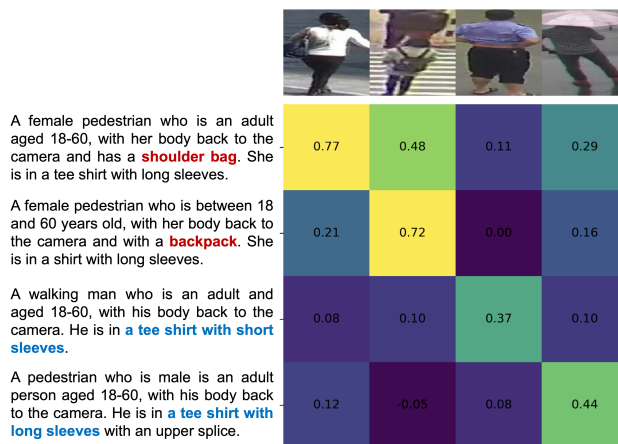

Figure 4. Visualization of similarity between different input texts and pedestrians with different attributes. Note that bright yellow represents high similarity, while dark blue represents low similarity.

huge backbone. Our vit-huge multi-task large model outperforms the vit-huge single-task large model and the vit-base large model, as can be observed. Meanwhile, it can achieve the state-of-the-art performance that is more advanced than baselines. Considering that large models need to occupy more hardware resources and take more training time, we provide Open-TransMind with vit-base backbone as our challenge code base.

5.2. Cross-Modal Image Retrieval

Implementation Details. In the experiment, we use a transformer-based vit-base model as the visual feature encoder in the cross-modal Open-TransMind, and a naive transformer as the text feature encoder. We use an AdamW optimizer with an initial learning rate of 0.0001 and weight decay of 0.0005, and we adopt piecewise decay to reduce the learning rate. We only use basic resizing for data augmentation. The entire training process is conducted on 8 A100 GPUs for 20 epochs with a batch size of 128 for each GPU.

Quantitative Results. We conduct four comparative experiments: (1) training from scratch; (2) directly loading CLIP pre-training weights; (3) fine-tuning on the basis of CLIP pre-training weights; and (4) loading CLIP pre-training weights for fine-tuning and using the All-in-One model in Track 1 in an off-the-shelf manner, as an additional visual feature encoder to provide richer semantic information about the traffic scene.

As shown in Table 3 of the experimental results, we can see that the CLIP pre-training weights already have good multi-traffic participant text and image retrieval capabilities. After fine-tuning with real traffic scene data, CLIP achieves even better results. Among all the comparative experiments,

Figure 5. Visualization of similarity between different input texts and vehicles with different attributes. Note that bright yellow represents high similarity, while dark blue represents low similarity.

the "with All in One off-the-shelf" method achieves the best results, which indicates that the All-in-One model has excellent traffic participant feature extraction ability through learning a large amount of traffic scene data.

Qualitative Results. Figure 4 and Figure 5 show the visualization of text-image retrieval, from which it can be seen that our model can effectively distinguish fine-grained differences in cross-modal information. In the pedestrian retrieval example in Figure 4, the query text in the first and second lines has the same attributes except for the different bag attributes. Open-TransMind can effectively distinguish this detailed attribute. The query text in the third and fourth lines has slight differences in the sleeve length of the shirt, and Open-TransMind can also provide correct retrieval results. In the vehicle retrieval example shown in Figure 5, the keywords in the query text in the first line are "blue", "Dongfeng brand", and "truck". Although the remaining three challenging query text changes the three attributes respectively, Open-TransMind can provide correct retrieval results. In addition, from this example, we can see that our model has certain knowledge emergence capabilities: there are three types of text and corresponding images in the training data: "Blue Mercedes Benz", "Dongfeng brand", and "Truck", but there is no text and corresponding images of "Blue Dongfeng Truck". The trained Open-TransMind can effectively complete the retrieval of "Blue Dongfeng Truck", as shown in the first column of row 1.

6. Conclusion

We propose a new baseline and benchmark for the first foundation model challenge of intelligent transportation, called Open-TransMind. Meanwhile, Open-TransMind is the first open-source foundation model of intelligent transportation with multi-task and multi-modal capabili-

ties that we are aware of. Furthermore, Open-TransMind can achieve cutting-edge performance on traffic scenario detection, classification, and segmentation datasets, while also possessing cross-modal text-image retrieval capabilities. Contestants and researchers can investigate new solutions for multi-task co-evolution as well as the interrelationships of cross-modal. In the future, we plan to add more tasks, such as text recognition and point cloud detection. More researchers are expected to join the investigation of foundation model technology.

References

[1] Ardhendu Behera, Zachary Wharton, Pradeep RPG Hewage, and Asish Bera. Context-aware attentional pooling (cap) for fine-grained visual classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 929–937, 2021. 5

[2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. 3

[3] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once for all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2020. 1

[4] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2

[5] Lisha Cui, Pei Lv, Xiaoheng Jiang, Zhimin Gao, Bing Zhou, Luming Zhang, Ling Shao, and Mingliang Xu. Context-aware block net for small object detection. *IEEE Transactions on cybernetics*, 52(4):2300–2313, 2020. 5

[6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. 3

[7] Zhen Dong, Yuwei Wu, Mingtao Pei, and Yunde Jia. Vehicle type classification using a semisupervised convolutional neural network. *IEEE transactions on intelligent transportation systems*, 16(4):2247–2256, 2015. 4

[8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 3

[9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012. 2

[10] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009, June 2022. 3

[11] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR, 18–24 Jul 2021. 3

[12] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 2, 3

[13] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Junjie Yan, and Xiaogang Wang. Hydraplus-net: Attentive deep features for pedestrian analysis. In *Proceedings of the IEEE international conference on computer vision*, pages 1–9, 2017. 2, 4

[14] Yuexin Ma, Xinge Zhu, Sibo Zhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6120–6127, 2019. 2

[15] S. Nag, S. Adak, and S. Das. What's there in the dark. In *26th IEEE International Conference on Image Processing (ICIP)*, 2019. 5

[16] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. 3

[17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 3

[18] Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. Hat: Hardware-aware transformers for efficient natural language processing. *ArXiv*, abs/2005.14187, 2020. 1

[19] Teng Xi, Yifan Sun, Deli Yu, Bi Li, Nan Peng, Gang Zhang, and et al. Ufo:unified feature optimization. In *European Conference on Computer Vision*, 2022. 1, 3, 5

[20] Xiaoqing Ye, Mao Shu, Hanyu Li, Yifeng Shi, Yingying Li, Guangjie Wang, Xiao Tan, and Errui Ding. Rope3d: the roadside perception dataset for autonomous driving and monocular 3d object detection task. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21341–21350, 2022. 2

[21] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2633–2642, 2018. 2, 3

[22] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, et al. Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21361–21370, 2022. 2

[23] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. M. Ni, and H. Y. Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv e-prints*, 2022. 5

[24] S. Zheng, J. Lu, H. Zhao, X. Zhu, Z. Luo, Y. Wang, Y. Fu, J. Feng, T. Xiang, and Phs Torr. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Computer Vision and Pattern Recognition*, 2021. 5

[25] Zhe Zhu, Dun Liang, Songhai Zhang, Xiaolei Huang, Baoli Li, and Shimin Hu. Traffic-sign detection and classification in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 3