

PAI-Diffusion: Constructing and Serving a Family of Open Chinese Diffusion Models for Text-to-image Synthesis on the Cloud

Chengyu Wang¹, Zhongjie Duan^{1,2}, Bingyan Liu^{1,3}, Xinyi Zou¹, Cen Chen², Kui Jia³, Jun Huang^{1*}

¹ Alibaba Group, Hangzhou, China

² East China Normal University, Shanghai, China

³ South China University of Technology, Guangzhou, China

{chengyu.wcy, zouxinyi.zxy, huangjun.hj}@alibaba-inc.com
eeliubingyan@mail.scut.edu.cn, zjduan@stu.ecnu.edu.cn
cenchen@dase.ecnu.edu.cn, kuijia@gmail.com

Abstract

Text-to-image synthesis for the Chinese language poses unique challenges due to its large vocabulary size, and intricate character relationships. While existing diffusion models have shown promise in generating images from textual descriptions, they often neglect domain-specific contexts and lack robustness in handling the Chinese language. This paper introduces *PAI-Diffusion*, a comprehensive framework that addresses these limitations. *PAI-Diffusion* incorporates both general and domain-specific Chinese diffusion models, enabling the generation of contextually relevant images. It explores the potential of using LoRA and ControlNet for fine-grained image style transfer and image editing, empowering users with enhanced control over image generation. Moreover, *PAI-Diffusion* seamlessly integrates with Alibaba Cloud’s Platform for AI, providing accessible and scalable solutions. All the Chinese diffusion model checkpoints, LoRAs, and ControlNets, including domain-specific ones, are publicly available. A user-friendly Chinese WebUI and the *diffusers-api* elastic inference toolkit, also open-sourced, further facilitate the easy deployment of *PAI-Diffusion* models in various local and cloud environments, making it a valuable resource for Chinese text-to-image synthesis.¹

1 Introduction

Recently, diffusion models (Rombach et al., 2022; Saharia et al., 2022) have emerged to address the challenges of generating realistic and high-quality images from textual descriptions. This research area has obtained widespread attention, driven by the increasing demand for automated image synthesis in various applications, such as art design, virtual reality, etc (Cao et al., 2023; Sun et al., 2023).

*Corresponding author.

¹Video presentation: https://atp-modelzoo-sh.oss-cn-shanghai.aliyuncs.com/release/conf_demo/acl_demo_2024.mov.

In the community, Stable Diffusion² has gained significant popularity due to its ability to generate high-quality images that align well with textual descriptions. However, when it comes to handling the Chinese language, similar models encounter certain challenges that hinder its performance. From the linguistic aspect, Chinese is a morphologically rich language with a large vocabulary size and complex inter-dependencies between characters. Furthermore, Chinese characters often have multiple meanings and can be combined to form compound words, making it challenging to establish accurate and consistent mappings between textual descriptions and visual representations (Liu et al., 2022).

Previously, in the literature, several works have been proposed to make specialized adaptations of diffusion models to improve the performance of text-to-image synthesis for Chinese (Wang et al., 2022c; Chen et al., 2023; Hu et al., 2023). Yet, there are still some notable drawbacks that need to be addressed. i) Many existing models focus on generating images based on generic textual descriptions, neglecting the ability to generate images in specific domains or contexts. ii) For the Chinese language, the potential of using LoRA (Hu et al., 2022) and ControlNet (Zhang and Agrawala, 2023) for fine-grained image style transfer and image editing have not been fully explored. iii) The lack of support for cloud-based product integration is another important drawback. Given the increasing popularity of cloud-based services and the demand for scalable and accessible text-to-image solutions, it is crucial to develop models and solutions that can be easily integrated into cloud platforms or deployed as cloud-based services.

In this work, we formally present *PAI-Diffusion*, which consists of a family of open Chinese diffusion models, together with user-friendly toolkits to serve these models on the cloud. Major features of *PAI-Diffusion* include the following:

²<https://stability.ai/stablediffusion>

- *PAI-Diffusion* incorporates both general and domain-specific Chinese diffusion models, specifically allowing for the generation of images that are tailored to specific contexts or domains (such as Chinese cuisine, poetry and paintings). This enables users to create visually compelling and relevant images for various domain-specific applications.
- *PAI-Diffusion* explores the potential of using LoRA and ControlNet for fine-grained style transfer and image editing, with a variety of corresponding Chinese models released. This empowers users with greater control over the generated images, enabling them to manipulate fine-grained semantic attributes and modify visual features based on their preferences.
- *PAI-Diffusion* extends our previous work (Liu et al., 2023) and ensures seamless integration with our Platform for AI (PAI) of Alibaba Cloud³, providing users with accessible and scalable solutions. By integrating with cloud products, *PAI-Diffusion* enables users to harness the power of cloud computing and enjoy the benefits of user-friendly and resource-efficient systems. For designers, our Chinese WebUI toolkit largely extends the Stable Diffusion WebUI⁴ to enrich its abilities to address the issues of the Chinese language. For application developers, our elastic inference toolkit *diffusers-api* makes it easy to deploy these models as RESTful web services. It further supports our compiler performance optimization tool named *PAI-Blade* (Zhu et al., 2021). When the functionality is enabled, the image generation speed is improved by 2-3 times compared to native PyTorch implementation, while maintaining the effectiveness.

To realize our promise for the *openness* of our research, we have taken the following actions to contribute *PAI-Diffusion* to the community:

- All the diffusion models, LoRAs and ControlNets of *PAI-Diffusion*, including domain-specific ones, have been released in our organizational Hugging Face repository⁵. Users can easily download and fine-tune all the models

³<https://www.alibabacloud.com/product/machine-learning>

⁴<https://github.com/AUTOMATIC1111/stable-diffusion-webui>

⁵<https://huggingface.co/alibaba-pai>

in order to support their own domain-specific applications.

- Our Chinese WebUI and *diffusers-api* toolkits have also been made publicly available⁶, so that *PAI-Diffusion* models are easy to be deployed in other environments beyond the Alibaba Cloud ecosystem, for both designers and application developers.
- All our models and codes are released under the Apache License (Version 2.0) to support both academic and commercial use.

2 Models

We introduce the models of *PAI-Diffusion*, a family of open Chinese diffusion models designed to address the challenges of generating high-quality images from Chinese textual descriptions.

2.1 Model Zoo

The model zoo of *PAI-Diffusion* consists of a collection of over ten open-source models, including base diffusion models, together with their corresponding LoRAs and ControlNets. A summary of these models is presented in Table 1, with a few cases of generated images presented in Figure 1. Readers can also find these models from our organizational Hugging Face repository described previously. Note that the list of models is not static. More new models will be added to the repository when they are ready to be released.

2.2 Model Architecture

To ensure full compatibility with the open-source community, our model architectures generally follow the *de facto* standard practice of Stable Diffusion (Rombach et al., 2022) where a CLIP-based text encoder, a U-Net and a VAE model are leveraged to obtain text embeddings, generate image embeddings in the latent diffusion space and decode the image for output, respectively. Particularly, the architectures of the U-Nets and VAEs of our large and xlarge diffusion models are in line with Stable Diffusion 1.5 and 2.1, respectively. Note that the difference between “large” and “xlarge” models is the size of generated images, rather than the number of parameters.

⁶The Chinese WebUI extension is released under the EasyNLP (Wang et al., 2022a,b) framework: https://github.com/alibaba/EasyNLP/tree/master/diffusion/chinese_sd_webui
The *diffusers-api* repository: <https://github.com/alibaba/diffusers-api>




Figure 1: Some examples of the generated images by *PAI-Diffusion* models. Prompts are originally in Chinese and have been manually translated into English for reference.

Model Name	#Parameters	Image Size (Default)	Domain
<i>pai-diffusion-general-large-zh</i>	1.04B	512×512	General purpose
-controlnet-canny	361M	512×512	General purpose
-controlnet-depth	361M	512×512	General purpose
<i>pai-diffusion-general-xlarge-zh</i>	1.04B	768×768	General purpose
<i>pai-diffusion-artist-large-zh</i>	1.04B	512×512	Artistic pictures
-controlnet-canny	361M	512×512	Artistic pictures
-controlnet-depth	361M	512×512	Artistic pictures
-lora-poem	25.5M	512×512	Paintings for Chinese poems
-lora-2.5d	25.5M	512×512	2.5D-style arts
<i>pai-diffusion-artist-xlarge-zh</i>	1.04B	768×768	Artistic pictures
<i>pai-diffusion-food-large-zh</i>	1.04B	512×512	Chinese cuisines
<i>pai-diffusion-anime-large-zh</i>	1.04B	768×512	Cartoon characters (anime)

Table 1: A summary of Chinese diffusion models, LoRAs and ControlNets released by us.

As for the Chinese language, we follow our previous work (Liu et al., 2023) to employ both 100 million text-image pairs from Wukong (Gu et al., 2022) and the largest Chinese KG available to us, i.e., OpenKG⁷ as our knowledge source to pre-train a Chinese knowledge-enhanced CLIP model that better understand the morphologically rich semantics of the Chinese language. The resources for training general-purpose models are also the same as in (Liu et al., 2023). For more details, we refer our readers to the original paper.

We have also collected a variety of domain-specific datasets to produce domain-specific diffusion models or LoRAs, depending on the volume of the corresponding datasets. Such domains include artistic pictures, paintings for Chinese poems, 2.5D-style arts, Chinese cuisines and cartoon characters. Note that we try our best to ensure that our models are built on legally and ethically sourced data. We specifically filter out any images that may have the probability to exhibit some degree of ethical bias. We obtain the permission to use in-house datasets (such as Chinese cuisine) to train and release the corresponding models.

ControlNet (Zhang and Agrawala, 2023) operates by incorporating a set of control vectors that encode specific image attributes or features. These control vectors are then incorporated into the residual blocks of the diffusion models. Furthermore, ControlNet allows for interactive image editing, enabling users to iteratively refine and modify the generated images, based on our WebUI toolkit. To enable fine-grained control, *PAI-Diffusion* integrates several ControlNets with Chinese diffusion models, facilitating the modification and customization of generated images based on user preferences. Currently, we have released the ControlNets based on canny edge detection (Canny, 1986) and Midas depth maps (Ranftl et al., 2022). Users can train their own ControlNets, using the same methods as ControlNet for Stable Diffusion.

2.3 Applications

PAI-Diffusion opens up a wide range of applications for the Chinese language. Here, we highlight some potential applications of *PAI-Diffusion*.

2.3.1 Artistic Creations and Design

Diffusion models revolutionize the way sketch images are transformed into captivating artworks.

With our models, users can witness their sketch images come to life based on the image-to-image pipeline. Examples can be found in Figure 2. We can see that our models empower artists to explore their creative boundaries and produce truly unique and mesmerizing artistic creations.

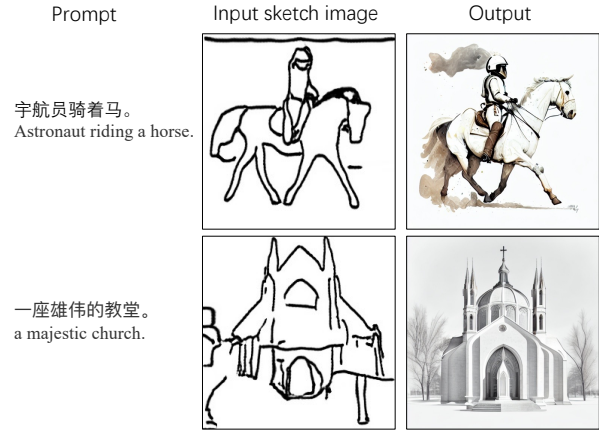


Figure 2: Two examples of artistic creations from sketch images using the image-to-image pipeline.

2.3.2 Cultural Preservation and Heritage

Our models present an elegant approach to restoring ancient Chinese paintings through image inpainting techniques. By harnessing the power of diffusion, these applications enable the recreation of missing or damaged areas in the paintings, seamlessly blending them with the original artwork (with examples shown in Figure 3). Through the application, ancient Chinese paintings once damaged or fragmented can be revitalized, allowing future generations to appreciate the cultural significance of invaluable artistic treasures.

2.3.3 Visual Reality Generation

Diffusion models offer a cutting-edge approach to creating immersive virtual reality experiences. Take traditional Chinese garden architecture as an example (see Figure 4). Our model blends various elements such as architecture, house and furniture to generate virtual scenes reminiscent of ancient Chinese gardens. Such applications serve as a valuable tool for people to experience and appreciate the beauty of ancient Chinese culture regardless of their physical location.

2.3.4 Others

Apart from the three examples, our models can be harnessed in other domains such as fashion design, interior decor, and even product development. For

⁷<http://openkg.cn/>




Figure 3: Two examples of the restoration of ancient Chinese paintings via image in-painting.

instance, our model can create harmonious and visually stunning environments for interior decor. In fashion design, our diffusion models can be applied to blend both traditional and contemporary styles, creating unique and captivating clothing designs. We do not further elaborate.

3 Toolkits

PAI-Diffusion provides user-friendly toolkits that facilitate the deployment and usage of diffusion models for Chinese text-to-image synthesis. In this section, we briefly introduce our two toolkits.

3.1 Chinese WebUI Toolkit

The Chinese WebUI toolkit is an extension to the Stable Diffusion WebUI specifically developed to support our Chinese models. It offers a web-based graphical interface that allows users (especially designers without programming expertise) to interact with the diffusion models easily. The toolkit provides a seamless user experience, enabling users to perform image synthesizing and editing. Snapshots of the toolkit are shown in Figure 6. Additionally, our toolkit offers two modes: *single-node* mode and *cluster* mode (shown in Figure 5), providing users with options based on their specific requirements and resources, introduced as follows:

- **Single-node:** It is suitable for scenarios when users want to quickly set up and use the toolkit



Figure 4: Four examples of scene generation for Chinese garden architecture using the text-to-image pipeline.

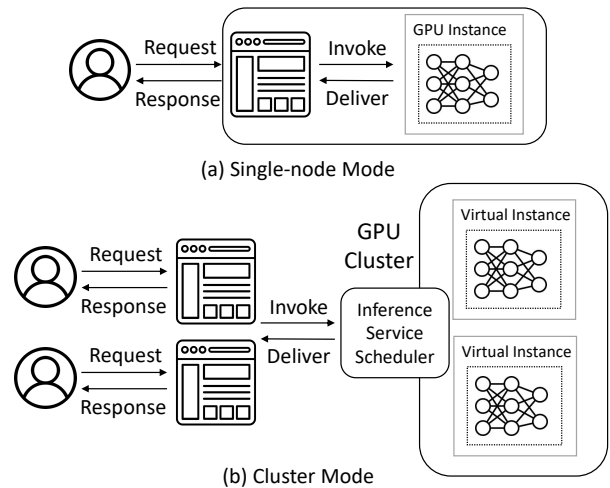


Figure 5: The system architectures of two modes to deploy our Chinese WebUI toolkit online.

with exclusive computational resources. It is particularly useful for individual users or small-scale deployments.

- **Cluster:** Leveraging the elastic inference service of PAI, the cluster is geared towards users who require high scalability and performance. In this mode, the toolkit utilizes a cluster of nodes to handle the workload, enabling parallel processing and efficient utilization of resources. It ensures efficient scaling and responsiveness, making it suitable for enterprise-level deployments or applications that require extensive computational power.

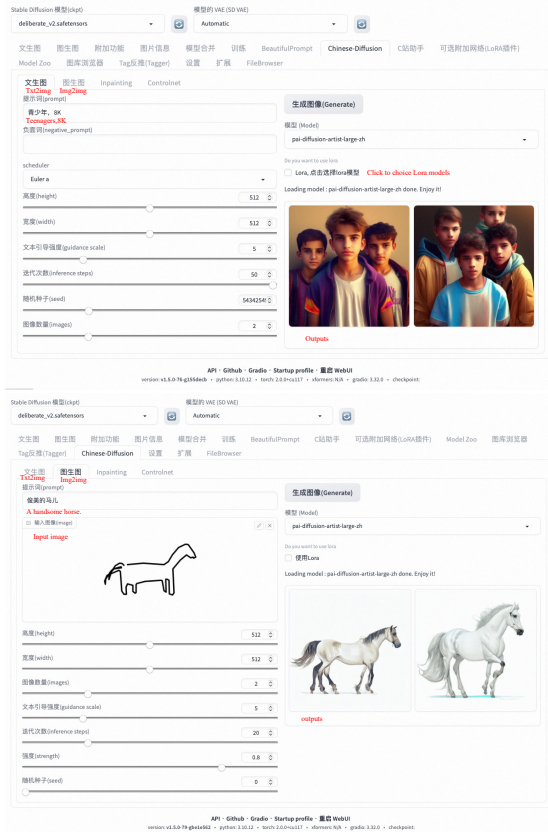


Figure 6: The snapshots of Chinese WebUI for text-to-image and image-to-image synthesis.

3.2 The *diffusers-api* Toolkit

The *diffusers-api* toolkit is built upon the Hugging Face *diffusers* library⁸ that provides a cloud service implementation for *PAI-Diffusion* models based on the PAI elastic inference service.⁹ Especially, we implement several inference pipelines customized for our Chinese models. Users can take advantage of a number of functions when sending HTTP requests, including text-to-image synthesis, image-to-image synthesis, image in-painting, image editing, etc. Below we show a sample request body for calling the *diffusers-api* service, with Chinese prompts manually translated into English for better understanding:

```
{
  "task_id": "001",
  "prompt": "romantic starry sky",
  "negative_prompt": "noise, low-quality",
  "func_name": "t2i",
  "steps": 25,
  "image_num": 1,
  "width": 512,
```

⁸<https://github.com/huggingface/diffusers>

⁹Note that *diffusers-api* is also compatible with original Stable Diffusion, which is not the focus of this paper.

Settings	PyTorch Native	<i>diffusers-api</i>
Inference time (s)	6.34	2.96
GPU memory (GB)	6.94	5.56

Table 2: The performance of *diffusers-api* for online deployment of our diffusion model.

```
"height": 512,
"use_base64": True
}
```

where “func_name” refers to the specific functionality that the user wish to use (“t2i” refers to text-to-image in this case). We release the source code of *diffusers-api* to provide users with the ability to quickly develop customized API services. This allows for easier implementation of any desired functionality, tailored to specific requirements. For example, users can build and customize their own schedulers (Duan et al., 2023).

In addition, *diffusers-api* leverages the AI compiler *PAI-Blade* (Zhu et al., 2021) for inference optimization. It significantly reduces the end-to-end latency of the inference processes and the GPU memory consumption, which ensures improved performance and efficiency in generating high-quality images, without any precision loss in computation. To verify the correctness of our argument, we have deployed our Chinese diffusion model (the large version) online using an NVIDIA A10 GPU with 50 sampling steps for inference. The generated image size is fixed to 512×512. We repeat the experiments in 20 times and report the averaged results in Table 2, which clearly prove the correctness of our claim.

4 Conclusion

This paper presents *PAI-Diffusion* to address the challenges in Chinese text-to-image synthesis. *PAI-Diffusion* integrates both general and domain-specific Chinese diffusion models, LoRAs and ControlNets, enabling the generation of contextually relevant images. Moreover, *PAI-Diffusion* ensures seamless integration with our cloud platform, offering accessible and scalable solutions. The release of models further encourages collaboration and innovation in the field. The public availability of Chinese WebUI and *diffusers-api* toolkits simplifies deployment in various environments. Our advancements pave the way for further research and development in Chinese text-to-image synthesis and relevant applications.

Acknowledgements

The authors would like to thank other members of the Platform for AI (PAI) team of Alibaba Cloud for the help of this work. This work was in part supported by the National Natural Science Foundation of China under grant number 62202170, Fundamental Research Funds for the Central Universities under grant number YBNLTS2023-014, Alibaba Group through Alibaba Innovative Research Program, and Alibaba Cloud through Research Talent Program with South China University of Technology.

Limitations

There are still some limitations that should be acknowledged. The effectiveness of our models heavily relies on accurate mappings between textual descriptions and visual representations. While efforts have been made to include a variety of domain-specific models, there may still be domains or contexts for which specialized models are not available. This limitation restricts the full potential in generating highly relevant and specific images for a wide range of applications. We suggest that users should further fine-tune our models if necessary.

Ethical Considerations

It is important to consider the ethical implications associated with its use and deployment. Diffusion models like *PAI-Diffusion* learn from large datasets, which may inadvertently contain biases present in the data. It is crucial to be aware of and mitigate any biases that may be perpetuated in the generated images. In addition, *PAI-Diffusion* has the potential to be misused for unethical purposes, such as generating inappropriate or harmful contents. Therefore, users should be encouraged to adhere to ethical standards and abide by terms and regulations when utilizing *PAI-Diffusion* models.

References

John F. Canny. 1986. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698.

Shidong Cao, Wenhao Chai, Shengyu Hao, Yanting Zhang, Hangyue Chen, and Gaoang Wang. 2023. Diffifashion: Reference-based fashion design with structure-aware transfer by diffusion models. *CoRR*, abs/2302.06826.

Zhongzhi Chen, Guang Liu, Bo-Wen Zhang, Qinghong Yang, and Ledell Wu. 2023. Altclip: Altering the language encoder in CLIP for extended language capabilities. In *ACL (Findings)*, pages 8666–8682. Association for Computational Linguistics.

Zhongjie Duan, Chengyu Wang, Cen Chen, Jun Huang, and Weining Qian. 2023. Optimal linear subspace search: Learning to construct fast and high-quality schedulers for diffusion models. *CoRR*, abs/2305.14677.

Jiayi Gu, Xiaojun Meng, Guansong Lu, Lu Hou, Minzhe Niu, Hang Xu, Xiaodan Liang, Wei Zhang, Xin Jiang, and Chunjing Xu. 2022. Wukong: 100 million large-scale chinese cross-modal pre-training dataset and A foundation framework. *CoRR*, abs/2202.06767.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. Lora: Low-rank adaptation of large language models. In *ICLR*. OpenReview.net.

Jinyi Hu, Xu Han, Xiaoyuan Yi, Yutong Chen, Wenhao Li, Zhiyuan Liu, and Maosong Sun. 2023. Efficient cross-lingual transfer for chinese stable diffusion with images as pivots. *CoRR*, abs/2305.11540.

Bingyan Liu, Weifeng Lin, Zhongjie Duan, Chengyu Wang, Ziheng Wu, Zhang Zipeng, Kui Jia, Lianwen Jin, Cen Chen, and Jun Huang. 2023. Rapid diffusion: Building domain-specific text-to-image synthesizers with fast inference speed. In *ACL*, pages 295–304. Association for Computational Linguistics.

Tingting Liu, Chengyu Wang, Xiangru Zhu, Lei Li, Minghui Qiu, Jun Huang, Ming Gao, and Yanghua Xiao. 2022. ARTIST: A transformer-based chinese text-to-image synthesizer digesting linguistic and world knowledge. In *EMNLP (Findings)*, pages 881–888. Association for Computational Linguistics.

René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. 2022. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(3):1623–1637.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10674–10685. IEEE.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. 2022. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*.

Shiqi Sun, Shancheng Fang, Qian He, and Wei Liu. 2023. Design booster: A text-guided diffusion model for image translation with spatial layout preservation. *CoRR*, abs/2302.02284.

- Chengyu Wang, Minghui Qiu, and Jun Huang. 2022a. Building natural language processing applications with easynlp. In *CIKM*, pages 5100–5101. ACM.
- Chengyu Wang, Minghui Qiu, Taolin Zhang, Tingting Liu, Lei Li, Jianing Wang, Ming Wang, Jun Huang, and Wei Lin. 2022b. Easynlp: A comprehensive and easy-to-use toolkit for natural language processing. In *EMNLP*, pages 22–29. Association for Computational Linguistics.
- Junjie Wang, Yuxiang Zhang, Lin Zhang, Ping Yang, Xinyu Gao, Ziwei Wu, Xiaoqun Dong, Junqing He, Jianheng Zhuo, Qi Yang, Yongfeng Huang, Xiayu Li, Yanghan Wu, Junyu Lu, Xinyu Zhu, Weifeng Chen, Ting Han, Kunhao Pan, Rui Wang, Hao Wang, Xiaojun Wu, Zhongshen Zeng, Chongpei Chen, Ruyi Gan, and Jiaying Zhang. 2022c. Fengshenbang 1.0: Being the foundation of chinese cognitive intelligence. *CoRR*, abs/2209.02970.
- Lvmin Zhang and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. *CoRR*, abs/2302.05543.
- Kai Zhu, Wenyi Zhao, Zhen Zheng, Tianyou Guo, Pengzhan Zhao, Junjie Bai, Jun Yang, Xiaoyong Liu, Lansong Diao, and Wei Lin. 2021. DISC: A dynamic shape compiler for machine learning workloads. In *EuroMLSys@EuroSys*, pages 89–95. ACM.