

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant 62472033, 92470203, U23A20314, 61972036), and the Beijing Natural Science Foundation (Grant L242022).

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L. Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikołaj Bińkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: A Visual Language Model for Few-Shot Learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 3
- [2] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 1
- [3] Xiaoyi Bao, Siyang Sun, Shuailei Ma, Kecheng Zheng, Yuxin Guo, Guosheng Zhao, Yun Zheng, and Xingang Wang. CoReS: Orchestrating the Dance of Reasoning and Segmentation. In *Computer Vision – ECCV 2024*, pages 187–204. Springer Nature Switzerland, Cham, 2025. 3, 6
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 1
- [5] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. COCO-Stuff: Thing and Stuff Classes in Context. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1209–1218, Salt Lake City, UT, USA, 2018. IEEE. 5, 1, 2
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *Computer Vision – ECCV 2018*, pages 833–851. Springer International Publishing, Cham, 2018. 2
- [7] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect What You Can: Detecting and Representing Objects Using Holistic Models and Body Parts. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1979–1986, Columbus, OH, USA, 2014. IEEE. 5, 1, 2
- [8] Yi-Chia Chen, Wei-Hua Li, Cheng Sun, Yu-Chiang Frank Wang, and Chu-Song Chen. Sam4mllm: Enhance multi-modal large language model for referring expression segmentation. In *European Conference on Computer Vision*, pages 323–340. Springer, 2024. 5, 6
- [9] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198, 2024. 1
- [10] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in neural information processing systems*, 34:17864–17875, 2021. 2
- [11] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention Mask Transformer for Universal Image Segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1280–1289, New Orleans, LA, USA, 2022. IEEE. 1, 2, 3
- [12] Yong Xien Chng, Henry Zheng, Yizeng Han, Xuchong Qiu, and Gao Huang. Mask grounding for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26573–26583, 2024. 3
- [13] 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, 2019. 3
- [14] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16321–16330, 2021. 3
- [15] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11583–11592, 2022. 2
- [16] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [17] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 7
- [18] Kai Fang, Anqi Zhang, Guangyu Gao, Jianbo Jiao, Chi Harold Liu, and Yunchao Wei. Combo: Conflict mitigation via branched optimization for class incremental segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25667–25676, 2025. 2
- [19] Guangyu Gao, Anqi Zhang, Jianbo Jiao, Chi Harold Liu, and Yunchao Wei. Prformer: Matching proposal and reference masks by semantic and spatial similarity for few-shot semantic segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025. 2
- [20] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer*

- vision and pattern recognition*, pages 6904–6913, 2017. 5, 1, 2
- [21] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 1
- [22] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617, 2018. 5, 1, 2
- [23] Kunyang Han, Yibo Hu, Mengxue Qu, Hailin Shi, Yao Zhao, and Yunchao Wei. ROSE: Revolutionizing Open-Set Dense Segmentation with Patch-Wise Perceptual Large Multimodal Model, 2025. 3
- [24] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2
- [25] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 5
- [26] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 5, 1, 2
- [27] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 2
- [28] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to Objects in Photographs of Natural Scenes. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 787–798, Doha, Qatar, 2014. Association for Computational Linguistics. 5, 6, 1, 2
- [29] Wonjae Kim, Bokyung Son, and Ildoo Kim. ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. 2
- [30] 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. Segment Anything. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3992–4003, Paris, France, 2023. IEEE. 1, 3, 5
- [31] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9579–9589, 2024. 1, 3, 5, 6, 2
- [32] Boyi Li, Kilian Q. Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven Semantic Segmentation, 2022. Comment: ICLR 2022. 2
- [33] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In *Advances in Neural Information Processing Systems*, pages 9694–9705. Curran Associates, Inc., 2021. 2
- [34] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. In *Proceedings of the 39th International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [35] Zhang Li, Biao Yang, Qiang Liu, Shuo Zhang, Zhiyin Ma, Liang Yin, Linger Deng, Yabo Sun, Yuliang Liu, and Xiang Bai. LIRA: Inferring Segmentation in Large Multimodal Models with Local Interleaved Region Assistance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 24056–24067, 2025. 3, 6
- [36] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7061–7070, 2023. 2
- [37] Chang Liu, Henghui Ding, and Xudong Jiang. GRES: Generalized Referring Expression Segmentation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23592–23601, Vancouver, BC, Canada, 2023. IEEE. 6
- [38] Chang Liu, Henghui Ding, and Xudong Jiang. Gres: Generalized referring expression segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23592–23601, 2023. 3
- [39] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual Instruction Tuning. *Advances in Neural Information Processing Systems*, 36:34892–34916, 2023. 1, 3, 5, 2
- [40] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 3
- [41] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [42] Timo Luddecke and Alexander Ecker. Image Segmentation Using Text and Image Prompts. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7076–7086, New Orleans, LA, USA, 2022. IEEE. 2
- [43] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019. 5, 1, 2
- [44] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and

- Alan Yuille. The Role of Context for Object Detection and Semantic Segmentation in the Wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, Columbus, OH, USA, 2014. IEEE. 7
- [45] Jie Qin, Jie Wu, Pengxiang Yan, Ming Li, Ren Yuxi, Xuefeng Xiao, Yitong Wang, Rui Wang, Shilei Wen, Xin Pan, et al. Freeseq: Unified, universal and open-vocabulary image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19446–19455, 2023. 2
- [46] 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 *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2
- [47] Vignesh Ramanathan, Anmol Kalia, Vladan Petrovic, Yi Wen, Baixue Zheng, Baishan Guo, Rui Wang, Aaron Marquez, Rama Kovvuri, Abhishek Kadian, Amir Mousavi, Yiwen Song, Abhimanyu Dubey, and Dhruv Mahajan. PACO: Parts and Attributes of Common Objects. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7141–7151, Vancouver, BC, Canada, 2023. IEEE. 5, 1, 2
- [48] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M. Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S. Khan. GLaMM: Pixel Grounding Large Multimodal Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13009–13018, 2024. 1, 3, 5
- [49] Zhongwei Ren, Zhicheng Huang, Yunchao Wei, Yao Zhao, Dongmei Fu, Jiashi Feng, and Xiaojie Jin. PixelLM: Pixel Reasoning with Large Multimodal Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26374–26383, 2024. 3, 5
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation, 2015. Comment: conditionally accepted at MIC-CAI 2015. 2
- [51] Tong Shao, Zhuotao Tian, Hang Zhao, and Jingyong Su. Explore the potential of clip for training-free open vocabulary semantic segmentation. In *European Conference on Computer Vision*, pages 139–156. Springer, 2024. 2
- [52] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. 5, 1, 2
- [53] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. FLAVA: A Foundational Language And Vision Alignment Model. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15617–15629, New Orleans, LA, USA, 2022. IEEE. 2
- [54] Lin Sun, Jiale Cao, Jin Xie, Xiaoheng Jiang, and Yanwei Pang. Cliper: Hierarchically improving spatial representation of clip for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23199–23209, 2025. 2
- [55] Hao Tang, Chenwei Xie, Haiyang Wang, Xiaoyi Bao, Tingyu Weng, Pandeng Li, Yun Zheng, and Liwei Wang. UFO: A Unified Approach to Fine-grained Visual Perception via Open-ended Language Interface, 2025. 2, 3, 5, 6
- [56] Jiajin Tang, Ge Zheng, Cheng Shi, and Sibeil Yang. Contrastive grouping with transformer for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23570–23580, 2023. 3
- [57] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1
- [58] Feng Wang, Jieru Mei, and Alan Yuille. Sclip: Rethinking self-attention for dense vision-language inference. In *European Conference on Computer Vision*, pages 315–332. Springer, 2024. 2
- [59] Hao Wang, Limeng Qiao, Zequn Jie, Zhijian Huang, Chengjian Feng, Qingfang Zheng, Lin Ma, Xiangyuan Lan, and Xiaodan Liang. X-SAM: From Segment Anything to Any Segmentation, 2025. Comment: Technical Report. 3
- [60] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution, 2024. Comment: Code is available at <https://github.com/QwenLM/Qwen2-VL>. *arXiv admin note: text overlap with arXiv:2408.15262* by other authors. 3
- [61] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 1, 2
- [62] Tao Wang, Changxu Cheng, Lingfeng Wang, Senda Chen, and Wuyue Zhao. HiMTok: Learning Hierarchical Mask Tokens for Image Segmentation with Large Multimodal Model, 2025. Comment: Accepted by ICCV 2025; the code is at <https://github.com/yayafengzi/LMM-HiMTok>. 3, 5, 6, 7
- [63] Cong Wei, Haoxian Tan, Yujie Zhong, Yujiu Yang, and Lin Ma. Lasagna: Language-based segmentation assistant for complex queries. *arXiv preprint arXiv:2404.08506*, 2024. 5, 6
- [64] Cong Wei, Yujie Zhong, Haoxian Tan, Yong Liu, Zheng Zhao, Jie Hu, and Yujiu Yang. HyperSeg: Towards Universal Visual Segmentation with Large Language Model, 2024. 3, 5, 7
- [65] Jiannan Wu, Muyan Zhong, Sen Xing, Zeqiang Lai, Zhaoyang Liu, Zhe Chen, Wenhui Wang, Xizhou Zhu, Lewei Lu, and Tong Lu. Visionllm v2: An end-to-end generalist multimodal large language model for hundreds of vision-language tasks. *Advances in Neural Information Processing Systems*, 37:69925–69975, 2024. 3, 5

- [66] Zhuofan Xia, Dongchen Han, Yizeng Han, Xuran Pan, Shiji Song, and Gao Huang. Gsva: Generalized segmentation via multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3858–3869, 2024. 3, 5, 6
- [67] Bin Xie, Jiale Cao, Jin Xie, Fahad Shahbaz Khan, and Yanwei Pang. Sed: A simple encoder-decoder for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3426–3436, 2024. 2
- [68] Jinjin Xu, Liwu Xu, Yuzhe Yang, Xiang Li, Fanyi Wang, Yanchun Xie, Yi-Jie Huang, and Yaqian Li. u-llava: Unifying multi-modal tasks via large language model. *arXiv preprint arXiv:2311.05348*, 2023. 5
- [69] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2945–2954, 2023. 2
- [70] Cilin Yan, Haochen Wang, Shilin Yan, Xiaolong Jiang, Yao Hu, Guoliang Kang, Weidi Xie, and Efstratios Gavves. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*, pages 98–115. Springer, 2024. 6
- [71] Zhao Yang, Jiaqi Wang, Yansong Tang, Kai Chen, Hengshuang Zhao, and Philip HS Torr. Lavt: Language-aware vision transformer for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18155–18165, 2022. 3
- [72] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. Modeling Context in Referring Expressions, 2016. 5, 6, 1, 2
- [73] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. *Advances in Neural Information Processing Systems*, 36:32215–32234, 2023. 2
- [74] Haobo Yuan, Xiangtai Li, Tao Zhang, Yueyi Sun, Zilong Huang, Shilin Xu, Shunping Ji, Yunhai Tong, Lu Qi, Jiashi Feng, and Ming-Hsuan Yang. Sa2VA: Marrying SAM2 with LLaVA for Dense Grounded Understanding of Images and Videos, 2025. Comment: Code: <https://github.com/Bytedance/Sa2VA>. 3
- [75] Anqi Zhang and Guangyu Gao. Background adaptation with residual modeling for exemplar-free class-incremental semantic segmentation. In *European Conference on Computer Vision*, pages 166–183. Springer, 2024. 2
- [76] Anqi Zhang, Guangyu Gao, Jianbo Jiao, Chi H Liu, and Yunchao Wei. Bridge the points: Graph-based few-shot segment anything semantically. *Advances in Neural Information Processing Systems*, 37:33232–33261, 2024. 2
- [77] Tao Zhang, Xiangtai Li, Hao Fei, Haobo Yuan, Shengqiong Wu, Shunping Ji, Chen Change Loy, and Shuicheng Yan. Omg-llava: Bridging image-level, object-level, pixel-level reasoning and understanding. *Advances in neural information processing systems*, 37:71737–71767, 2024. 1, 3, 5, 6
- [78] Yichi Zhang, Ziqiao Ma, Xiaofeng Gao, Suhaila Shakiah, Qiaozi Gao, and Joyce Chai. Groundhog Grounding Large Language Models to Holistic Segmentation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14227–14238, Seattle, WA, USA, 2024. IEEE. 5
- [79] Zheng Zhang, Yeyao Ma, Enming Zhang, and Xiang Bai. PSALM: Pixelwise SegmentAtion with Large Multi-modal Model. In *Computer Vision – ECCV 2024*, pages 74–91. Springer Nature Switzerland, Cham, 2025. 3, 5, 6, 7
- [80] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic Understanding of Scenes through the ADE20K Dataset, 2018. 5, 7, 1, 2
- [81] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European conference on computer vision*, pages 696–712. Springer, 2022. 2
- [82] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models, 2023. Comment: Project Website: <https://minigpt-4.github.io/>; Code, Pretrained Model, and Dataset: <https://github.com/Vision-CAIR/MiniGPT-4>; Deyao Zhu and Jun Chen contributed equally to this work. 3
- [83] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025. 2, 3, 5

Rethinking MLLM Itself as a Segmenter with a Single Segmentation Token

Supplementary Material

6. Limitations

Our method achieves competitive results on various segmentation tasks, yet the limitations still exist. The token interactions among the [IMG] tokens and [SEG] token enhance the spatial precision of features, yet the redesigned attention mask becomes an obstacle for autoregressive inference and multi-round reasoning. We could only predefine the text templates or separate the process of text inference and segmentation as compromise. Besides, the original VQA capabilities of MLLMs are not fully preserved according to Sec. 8.1. The enhanced localization and grounding capabilities from segmentation samples conflict with OCR-oriented and more complex knowledge reasoning scenarios, presenting a promising direction for future work toward better balance.

7. Details about training datasets

We utilize a broad collection of vision–language and pixel-level segmentation datasets to train both the base version of SELFIE and the segmentation-enhanced SELFIE-SEG. The details are shown in Tab. 8. The VQA component is constructed from six datasets, where VQAv2 [20] provides large-scale human-annotated question–answer pairs for general vision understanding, LLaVA-150k [39] offers high-quality multimodal conversational annotations from GPT-4, and OKVQA [43], TextVQA [52], VizWiz [22], and GQA [26] further contribute knowledge-based, text-centric, low-quality-image, and compositional reasoning supervision. These datasets are incorporated without magnification ($1\times$) for SEG version, totaling 421k samples. For language-guided referring expression segmentation, we adopt the RefCOCO, RefCOCO+, and RefCOCOg datasets [28, 72], which feature object-level referring expressions with increasing linguistic complexity. These datasets are expanded by a $20\times$ SEG-rate, providing 1.12M effective samples for SELFIE-SEG. To strengthen dense pixel-level perception, we employ ADE20K [80] that covers a broad spectrum of indoor/outdoor scenes with fine-grained masks, along with COCO-Stuff [5] and Pascal-Part [7] for diverse semantic regions and part-level annotations, and LVIS-PACO [47], which supplies long-tailed, instance-rich perceptual concepts. Each dataset is magnified $6\times$, yielding 504k samples for SEG version. Finally, ReasonSeg [31] is included to support more complex reasoning-driven segmentation, where its limited 239 samples are expanded $6\times$ into approximately 1.4k effective instances. Overall, our training corpus comprises roughly 561k samples for the base version SELFIE and around 2.4M magnified samples for SELFIE-

SEG.

8. Additional Experiment Results

8.1. Experiment results of VQA

Across 2B and 8B model scales, SELFIE exhibits a consistent performance pattern when compared with the corresponding InternVL3 baselines, as illustrated in Tab. 9. On generic benchmarks such as VizWiz, GQA, and VQAv2, SELFIE consistently achieves similar results as InternVL3, especially with gains of 11.2% on 2B and 6.5% on 8B on VizWiz and moderate improvements on GQA. These results suggest that introducing segmentation-aware visual supervision could retain the original generic understanding ability of the images. By contrast, SELFIE shows lower performance on OKVQA and TextVQA at both scales. Since these benchmarks heavily depend on external knowledge grounding (OKVQA) or OCR-oriented textual reasoning (TextVQA), the performance gap indicates that segmentation-focused training, provides limited improvement in text-heavy or knowledge-intensive settings even with specific training data. A similar trend is shown on instruction-oriented multimodal benchmarks (MMB-en/cn, MME), where SELFIE trails InternVL3 regardless of scale. Most of the performance reduction is on the OCR-oriented sub-tasks and more complex knowledge sub-tasks. Nevertheless, SELFIE maintains competitive POPE scores across scales, matching or approaching InternVL3, demonstrating that stronger spatial grounding introduced by segmentation has a limited negative influence on hallucination.

Overall, the unified comparison across both 2B and 8B models demonstrates that the strengths of SELFIE lie primarily in perception robustness and grounding-oriented reasoning, enabled by segmentation-enhanced visual modeling, whereas performance trade-offs emerge on OCR-heavy and knowledge-driven benchmarks. This consistent pattern across scales highlights the complementary nature of segmentation-aware learning within MLLMs and reveals clear future directions for balancing visual grounding with textual and knowledge-centric capabilities.

8.2. Experiment results with other MLLMs

We conduct several experiments of SELFIE based on different versions of MLLMs in Tab. 10, including InternVL2-2B and InternVL2.5-2B, on which some of the previous methods applied. The results show that even with previous versions of InternVL, our approach still achieves competitive performance. Using InternVL2.5-2B as the base model attains similar results on the standard version of

Task	Dataset	Samples	SEG-rates	SEG-samples
VQA	VQAv2 [20]	100k	1×	100k
	OKVQA [43]	9k		9k
	TextVQA [52]	35k		35k
	VizWiz [22]	20k		20k
	GQA [26]	100k		100k
	LLaVA-150k [39]	157k		157k
Referring Expression Segmentation	RefCOCO [28]	17k	20×	340k
	RefCOCO+ [28]	17k		340k
	RefCOCog [72]	22k		440k
Semantic Segmentation	ADE20k [80]	20k	6×	120k
	COCOStuff [5]	30k		180k
	Pascal-Part [7]	4k		24k
	LVIS-PACO [47]	30k		180k
Reasoning Segmentation	ReasonSeg [31]	239	6×	1.4k
Overall		561k		2.4M

Table 8. Details of the multiple datasets for training. SEG-rates represent the magnification of the dataset samples for the training of SELF1E-SEG version.

Methods	VQAv2	OKVQA	VizWiz	GQA	TextVQA	POPE	MMB-en	MMB-cn	MME
InternVL3-2B [83]	80.1	56.1	56.6	60.7	77.0	89.6	81.1	78.4	2221.2
InternVL3-8B [83]	81.8	61.9	63.4	63.2	78.9	91.1	83.4	82.2	2415.4
SELF1E-2B	77.7	48.5	67.8	61.5	71.7	89.4	69.6	66.2	2014.6
SELF1E-8B	80.2	54.0	69.9	64.1	72.3	91.1	74.6	72.0	2265.1

Table 9. Comparison of the VQA performance of our approach with their original base MLLMs.

Method	MLLM	RefCOCO			RefCOCO+			RefCOCog	
		val	testA	testB	val	testA	testB	val	test
SELF1E-2B	InternVL3-2B	80.2	82.1	77.6	74.6	79.1	69.2	77.0	77.8
	InternVL2-2B	77.7	80.6	74.6	71.5	76.7	66.9	74.3	74.7
	InternVL2.5-2B	80.1	82.2	78.0	74.7	78.7	69.8	76.5	77.6
SELF1E-SEG-2B	InternVL3-2B	84.3	85.4	82.3	78.9	83.5	75.1	80.4	81.0
	InternVL2-2B	83.5	85.5	81.4	77.7	82.0	72.9	79.5	80.1
	InternVL2.5-2B	85.2	86.7	83.5	79.9	83.4	75.2	81.0	82.3

Table 10. Comparison of results with different MLLMs as base model on the Referring Expression Segmentation benchmarks (RefCOCO+/g).

SELF1E, while having approximately 1% of advantage on the SELF1E-SEG version. In summary, the state-of-the-art performance of our SELF1E does not heavily rely on the upgrade of MLLM, where it still has advanced performance with earlier versions of InternVL.

8.3. Ablation Study of Retaining Resolution

The RFR operation in Sec. 3.3 requires uncompressed image features for retaining the original resolution of image features from the encoder. Thus, we design an experiment to measure the effectiveness of different strategies. As shown in Tab. 11, we compare the self-replication strategy with the scanning strategy. To be specific, the original pixel-shuffle process has the same stride value as the fac-

	RefCOCO	RefCOCO+	RefCOCog
Baseline	76.2	72.3	74.5
Scanning	78.3	73.3	76.3
Self-Replication	78.9	73.5	76.3

Table 11. Ablation study on the operations for retaining the resolution.

tor, so that different groups of features are not overlapped. The scanning strategy set the stride as 1, which preserves the original resolution. However, the results with scanning strategy, although higher than the baseline without any strategy, are still slightly lower than self-replication 0.6% on RefCOCO and 0.2% on RefCOCO+. The single pixel features from the scanning strategy are generated from α pixels from pre-compressed image features, while those from the self-replication strategy only correspond to the same pixel that preserves more precise spatial details.

8.4. Efficiency Comparison

To clarify when decoder-free is preferable, we report inference efficiency in Tab. 12 based on a single NVIDIA RTX4090. SELF1E achieves the fastest inference, outperforming LISA with specialist segmentation decoder and UFO with multi-token prediction. Even without customization for higher efficiency, SELF1E is still more memory-efficient than LISA and significantly faster than UFO, as it eliminates computational efforts on auxiliary mask decoders and multi-token decoding.

Table 12. Efficiency comparison among LISA, UFO, and SELF1E.

Method	Inference (ms)	FPS	Memory (GB)
LISA-7B	250.0	4.00	19.2
UFO-7B	961.5	1.04	14.4
SELF1E-8B	105.0	9.52	17.7

8.5. More Analysis on VQA

As shown in Tab. 13, incorporating pixel-level supervision leads to moderate degradation on knowledge-intensive and OCR-related tasks (e.g., Artworks, OCR, Commonsense, Code), while improving spatial understanding (Position). This indicates that segmentation supervision biases the model toward spatial grounding at some cost to abstract reasoning. Representative VQA examples in Fig. 5 further illustrate this effect, where spatial and positional queries improve while format-sensitive or multi-step reasoning may degrade. Overall, SELF1E preserves general VLM capability reasonably well while making an explicit and transparent trade-off to enable high-quality segmentation.

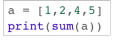


	The image shows a python code. Is the output of the code '5'? Please answer yes or no.	The image shows a python code. What is the output of the code?
InternVL3-2B	✔ no	✔ The output of the code is '12'.
SELF1E-2B	✘ yes	✔ The code creates a list of numbers, adds them together, and prints the result. The result is 12.
	I am going to play basketball. Is the shoe in the picture an appropriate choice?"	
InternVL3-2B	✔ The shoe in the picture is a classic dress shoe, which is more suitable for formal or business settings rather than playing basketball. For basketball, you would want a sneaker with better grip and support.	
SELF1E-2B	✘ Yes, the black leather shoe in the picture is an appropriate choice for playing basketball. It is a good-looking and stylish option that can complement a casual or sporty outfit. The shoe's design, with its leather material and laces, adds a touch of ...	
	Is the person under the kite?	
InternVL3-2B	✘ No, the person is not under the kite. They are standing on the beach, holding onto the kite's control lines while preparing to engage in kiteboarding or another water sport activity ...	
SELF1E-2B	✔ Yes, the person is under the kite while standing on the beach.	

Figure 5. Representative VQA results comparison.

9. Prompt Settings for Segmentation Tasks

9.1. Training

Our templates inherit the design principle of LISA [31]. Different dataset types use different prompt templates during training to align with their annotation styles.

For semantic segmentation task and vanilla referring expression segmentation task, we refer to the category name or object description simply as *text* for convenience:

We define a short question list:

- Can you segment the $\{text\}$ in this image?
- Please segment the $\{text\}$ in this image.
- What is $\{text\}$ in this image? Please respond with segmentation mask.
- What is $\{text\}$ in this image? Please output segmentation mask.

and an answer list:

- It is [SEG].
- Sure, [SEG].
- Sure, it is [SEG].
- Sure, the segmentation result is [SEG].
- [SEG].

The full template can be represent as:

USER: {a random choice from short question list}

ASSISTANT: {a random choice from answer list}

For reasoning segmentation task, the query expands into a longer, implicit instruction:

We use a long question list as below when the instruction is a full sentence; otherwise, we apply the short question list:

- $\{instruction\}$ Please respond with

Table 13. MME Benchmark Performance Comparison

	Method	Exist.	Count	Pos.	Color	Post.	Celeb.	Scene	Landm.	Artw.	OCR
Perp.	SELFIE	190.00	140.00	153.33	175.00	164.29	159.41	160.25	156.50	140.50	102.50
	InternVL3-2B	195.00	165.00	136.67	170.00	157.48	155.00	156.25	163.75	156.75	155.00
	Method	Comm.	Num.	Trans.	Code	Total Score					
Cog.	SELFIE	97.86	100.00	162.50	112.50	2014.64					
	InternVL3-2B	115.71	105.00	185.00	147.50	2164.11					

segmentation mask.

- `{instruction}` Please output segmentation mask.

The full template can be represent as:

USER: {a random choice from long or short question list}

ASSISTANT: {a random choice from answer list}

9.2. Validation

Our validation prompts follow two instruction formats, depending on whether the dataset provides object names or full-sentence instructions. Below, we provide the exact input–output templates used during validation.

For giving a specific object or description (*i.e.* RES and OVS datasets):

USER: What is {object’s name or description} in this image? Please output segmentation mask.

ASSISTANT: [SEG].

For giving a full sentence as instruction (*i.e.* ReasonSeg datasets):

USER: {Instruction} Please output segmentation mask.

ASSISTANT: [SEG].

10. Additional Visualization Results

10.1. Reasoning Segmentation

Reasoning segmentation requires model to infer the correct target object from implicit instructions, rather than relying on explicit object names. As shown in Fig. 6, SELFIE demonstrates strong capability in interpreting complex linguistic instructions and localizing the correct regions with high spatial precision. Although our architectural modifications primarily focus on visual features, the LLM’s reasoning ability remains unaffected, retaining its powerful linguistic inference capacity. Furthermore, the increase in the mask’s native resolution provides more detailed structural cues, enabling the model to better capture fine object boundaries. Overall, these results confirm that SELFIE maintains strong reasoning capabilities while benefiting from enhanced visual precision, leading to accurate segmentation under complex reasoning instructions.

10.2. Open-Vocabulary Segmentation

Fig. 7 presents the visualization results for open-vocabulary segmentation. It is important to note that masks with the same color across different images do not represent the same category; they are merely used to distinguish different objects within a single image. The main challenges in OVS lie in segmenting all instances of a category within an image and accurately distinguishing objects at boundaries, especially when multiple semantically similar objects are present. From our visualizations, SELFIE performs robustly even in images containing many categories. The model demonstrates precise classification, effectively distinguishing semantically similar objects and accurately capturing object boundaries. These results highlight the model’s strong generalization and fine-grained segmentation capabilities in complex, multi-category scenarios.

10.3. Token Interaction

We visualize the effects of different token interaction strategies on the RES task. Figure 8 presents the attention maps from the [SEG] token to [IMG] tokens. When only [IMG] to [IMG] attention is applied, the model is unable to access segmentation-related semantic cues from the [SEG] token. As a consequence, the attention maps sometimes fail to distinguish objects with similar semantics, particularly when the instruction specifies one target among multiple semantically related objects. This limitation becomes even more pronounced for location-dependent queries, where the model may incorrectly allocate high attention to a semantically similar but spatially incorrect object. These observations demonstrate the effectiveness of our [IMG] → [SEG] token interaction strategy, which substantially alleviates the issues discussed above.

	Instruction	Pred	GT
	When traveling by train, we often encounter obstacles such as rivers or valleys. What infrastructure in the picture helps the train cross over these obstacles? Please output segmentation mask.		
	What is a car with a color that is closer to lipstick color in this image? Please output segmentation mask.		
	What is the food that makes people feel spicy or hot in this image? Please output segmentation mask.		
	What is the more advanced aircraft in this image? Please output segmentation mask.		
	What part in the living room can people sit on and watch TV or take a nap? Please output segmentation mask.		
	When taking pictures with a camera, what part of the camera is responsible for focusing the image and adjusting the depth of field? Please output segmentation mask.		
	Many children enjoy building and creating things using small colorful blocks. What objects in the picture could be used to create various structures and shapes? Please output segmentation mask.		

Figure 6. Visualization results on ReasonSeg demonstrate outstanding reasoning ability of SELFIE. “Pred” denotes the predictions from SELFIE, “GT” denotes the ground-truth masks.




Figure 7. Visualization results on open-vocabulary segmentation. “OG” refers to the original image without overlays.

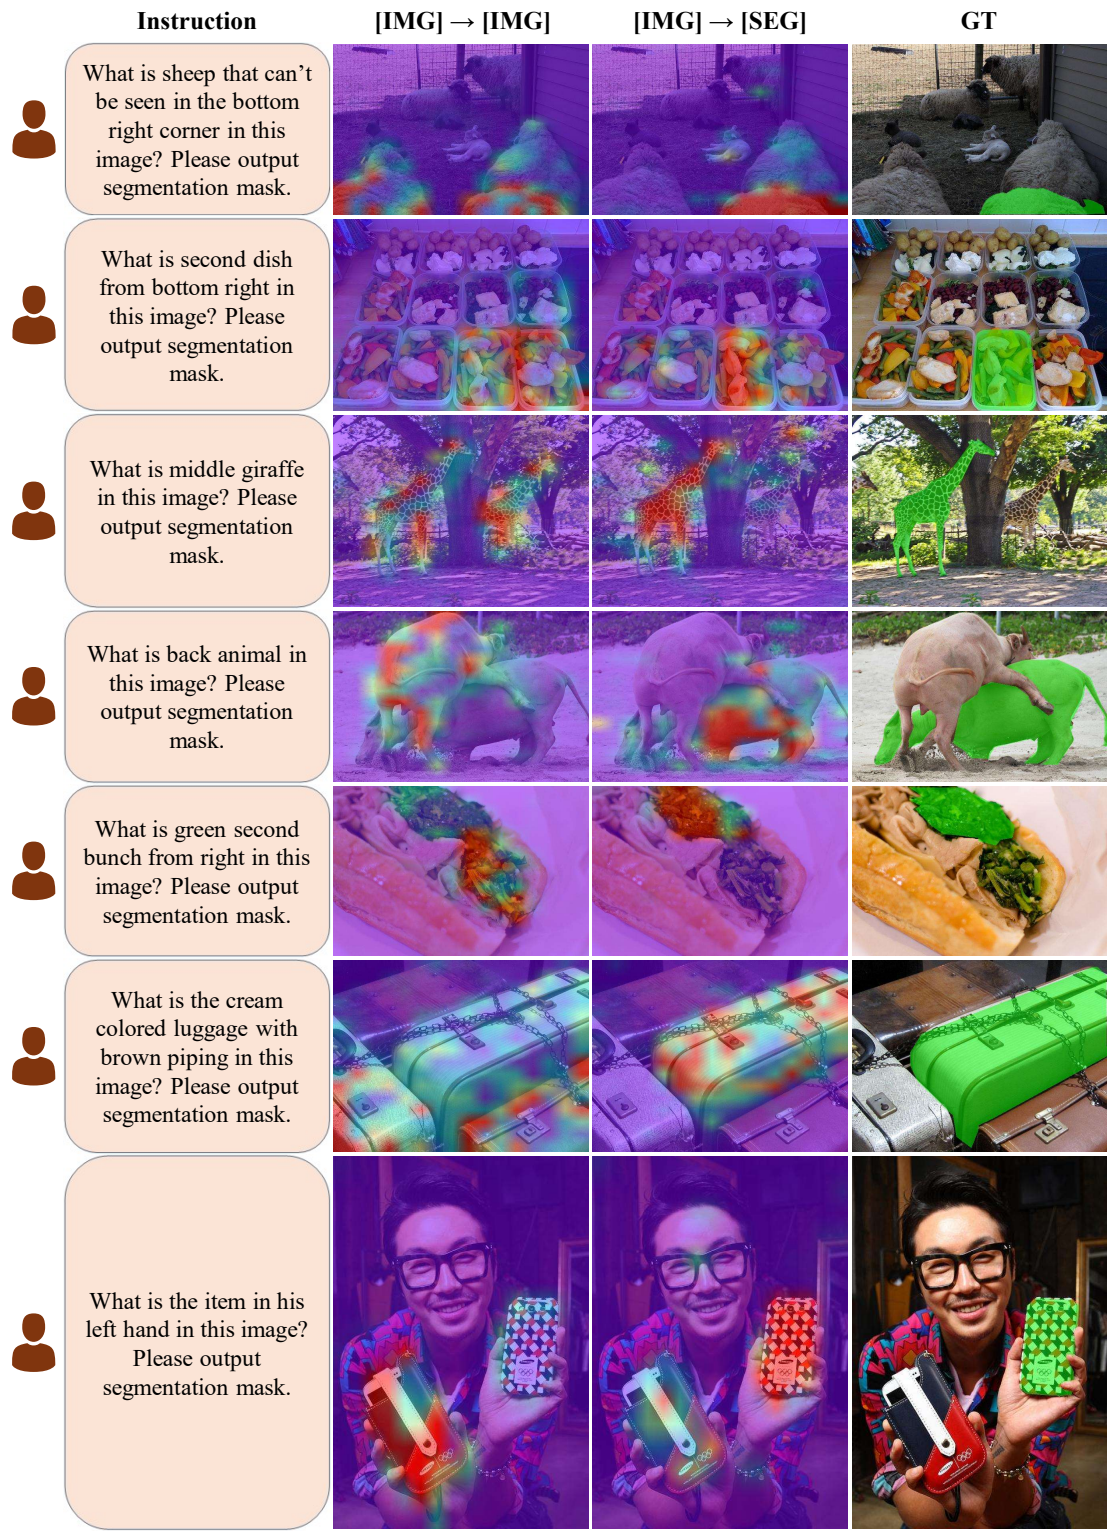


Figure 8. Visualization results show the attention maps of [SEG] to [IMG] tokens under different attention-mask designs. “[IMG] → [IMG]” indicates that all image tokens use a bidirectional attention mask, while all other tokens follow a causal mask. “[IMG] → [SEG]” means that, in addition to the bidirectional mask among image tokens, all image tokens are also allowed to interact with the [SEG] token.