







**Table 1: Important parameters that can be customized.**

Parameter	Description	Default Value
GPU_COUNT	The number of GPUs.	1
IMAGES_PER_GPU	The number of images to train with on each GPU.	1
STEPS_PER_EPOCH	The number of training steps per epoch.	1000
NUM_CLASSES	The number of classification classes (including background).	134
LEARNING_RATE	Learning rate.	0.001
LEARNING_MOMENTUM	Learning momentum.	0.9
IMAGE_PATH	The path of the images related files.	../data/
JSON_PATH	The path of the json files.	data
WEIGHT_PATH	The path of the default model weights.	models/CIN_ooi_all.pth
IMAGE_SIZE	The size of image after resizing and padding.	1024
MAP_IOU	The iou threshold when mapping prediction to the ground truth.	0.5
STUFF_THRESHOLD	The threshold when filtering small stuff.	1000
THING_NUM_CLASSES	The number of things (including background).	81
STUFF_NUM_CLASSES	The number of stuff.	53
SELECTION_THRESHOLD	The threshold when selecting IOIs.	0.4

```
<panoptic segmentation path> --sem_ext
<semantic segmentation path> --p_intr_ext
<interest estimation path> --sel_ext
<IOI selection method> --config
<configuration file path>
```

### 3 REPRODUCIBILITY EFFORTS

In the original paper, the connections among different files were not clearly described, and when the reviewers tried to run the scripts for the experiments, they found that some of the software dependencies such as “PyYAML” and “cffi” were missing in the given list and the “import” sentences were not updated in the original code. The authors quickly fixed the problems by offering a detailed explanation at the end of the Section 1.3. The missing dependencies also have been added in the Section 2.1 and the README file in the code repository has been updated. Moreover, a docker image containing all the software dependencies has been provided in response to the reviewers’ advice, making it more feasible for researchers in the area.

The reviewers acknowledge the efforts of the original authors to provide necessary corrections during the reviewing process, including simplifying the format of scripts parameters and fixing minor bugs in the code, as well as careful proofreading.

In conclusion, two reviewers and the authors worked together for this companion paper. The revised code now enables third-party researchers to reproduce the experiments in the original paper and is customizable for further research on visual relation detection.

### 4 CONCLUSION

In this paper, we provided the details of the artifacts of the paper “Instance of Interest Detection” for replication. The artifacts contain the IOID dataset and the source code for

experiments in the paper. Taking advantage of the source code, the experiments can be operated and customized.

### ACKNOWLEDGEMENT

This work is supported by Natural Science Foundation of Jiangsu Province (BK20191248), National Science Foundation of China (61732007), Science, Technology and Innovation Commission of Shenzhen Municipality (JCYJ20180307151516166), and Collaborative Innovation Center of Novel Software Technology and Industrialization.

### REFERENCES

- [1] Liangchieh Chen, George Papandreou, Iasonas Kokkinos, Kevin P Murphy, and Alan L Yuille. 2018. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets and Atrous Convolution and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018), 843–848.
- [2] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross B Girshick. 2017. Mask R-CNN. *IEEE Conference on Computer Vision and Pattern Recognition*.
- [3] Qibin Hou, Mingming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, and Philip H S Torr. 2017. Deeply Supervised Salient Object Detection with Short Connections. *IEEE Conference on Computer Vision and Pattern Recognition*.
- [4] Guanbin Li, Yuan Xie, Liang Lin, and Yizhou Yu. 2017. Instance-Level Salient Object Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [5] Nian Liu, Junwei Han, and Mingshan Yang. 2018. PiCANet: Learning Pixel-Wise Contextual Attention for Saliency Detection. *IEEE Conference on Computer Vision and Pattern Recognition*.
- [6] Zhiming Luo, Akshaya Kumar Mishra, Andrew Achkar, Justin A Eichel, Shaozi Li, and Pierre-marc Jodoin. 2017. Non-local Deep Features for Salient Object Detection. *IEEE Conference on Computer Vision and Pattern Recognition*.
- [7] Junting Pan, Cristian Cantonferrer, Kevin Mcguinness, Noel E Oconnor, Jordi Torres, Elisa Sayrol, and Xavier Giro I Nieto. 2017. SalGAN: Visual Saliency Prediction with Generative Adversarial Networks. *IEEE Conference on Computer Vision and Pattern Recognition*.
- [8] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. *International Conference on Machine Learning*.
- [9] Fan Yu, Haonan Wang, Tongwei Ren, Jinhui Tang, and Gangshan Wu. 2019. Instance of Interest Detection. *ACM International Conference on Multimedia*.