

A SELF-SUPERVISED METHOD FOR BODY PART SEGMENTATION AND KEYPOINT DETECTION OF RAT IMAGES

Áron Fóthi, László Kopácsi and András Lőrincz
(Budapest, Hungary)

Communicated by Péter Burcsi

(Received January 9, 2024; accepted April 26, 2024)

Abstract. Recognition of individual components and keypoint detection supported by instance segmentation is crucial to analyze the behavior of agents on the scene. Such systems could be used for surveillance, self-driving cars, and also for medical research, where behavior analysis of laboratory animals is used to confirm the aftereffects of a given medicine. A method capable of solving the aforementioned tasks usually requires a large amount of high-quality hand-annotated data, which takes time and money to produce. In this paper, we propose a method that alleviates the need for manual labeling of laboratory rats. To do so, first, we generate initial annotations with a computer vision-based approach, then through extensive augmentation, we train a deep neural network on the generated data. The final system is capable of instance segmentation, keypoint detection, and body part segmentation even when the objects are heavily occluded.

Key words and phrases: Computer vision, deep learning, keypoint detection, body part segmentation, instance segmentation.

2010 Mathematics Subject Classification: 68-06.

The research has been supported by the European Institute of Innovation and Technology. We want to thank Árpád Dobolyi and Dávid Keller for providing the database. Áron Fóthi and András Lőrincz were supported by the ELTE Institutional Excellence Program of the National Research, Development and Innovation Office (NKFIH-1157-8/2019-DT), and by the “Application Domain Specific Highly Reliable IT Solutions” project implemented with the support of the National Research, Development and Innovation Fund of Hungary under the Thematic Excellence Programme no. 2020-4.1.1.-TKP2020 (National Challenges Sub-programme) funding scheme, respectively.

1. Introduction

Body part segmentation, keypoint detection, and instance segmentation are critical for understanding interactions between agents on the scene. To address these tasks, we have to tackle the difficulties arising from heavy occlusions of objects. Moreover, in the case of behavior prediction of laboratory animals, such as rats, objects can be highly similar as well.

There are several deep neural network-based architectures [3, 16, 5] capable of solving these tasks and handling the challenges posed by the dataset. However, they all require a large set of carefully annotated samples. Although there are solutions that can reduce the need for massive databases and make the annotation process faster [11, 15], they still require some form of manual labeling.

In this paper, we propose a method that is able to automatically annotate rats without any human effort. To realize this, an initial segmentation of objects is generated via foreground-background segmentation, from which keypoints and parts of samples are separated using various computer vision (CV) techniques. Given the derived data and their corresponding CV generated labels, deep neural networks are exploited for tracking. We use the Mask R-CNN [5] architecture both for body part segmentation and for keypoint detection (including the instance segmentation) tasks. We study different augmentation techniques and handle occlusions as suggested in the video segmentation literature [8, 13].

We evaluate the final method on hand-annotated samples using metrics introduced in the COCO benchmark [9]. We started with an initial average precision (AP) of 53.22% on instance segmentation, 48.91% on keypoint detection, and 9.38% on body part segmentation, and by training deep models in a self-supervised manner, we achieved 61.92%, 77.53%, and 28.87%, respectively.

Our contributions are listed below:

- We propose a computer vision-based method for automatically annotating keypoints and body parts from foreground-background segmentation.
- We study the contributions of various augmentation methods used in the video segmentation literature.
- Finally, we train deep models on the generated labels in a self-supervised manner to handle the heavy occlusions present in the database.

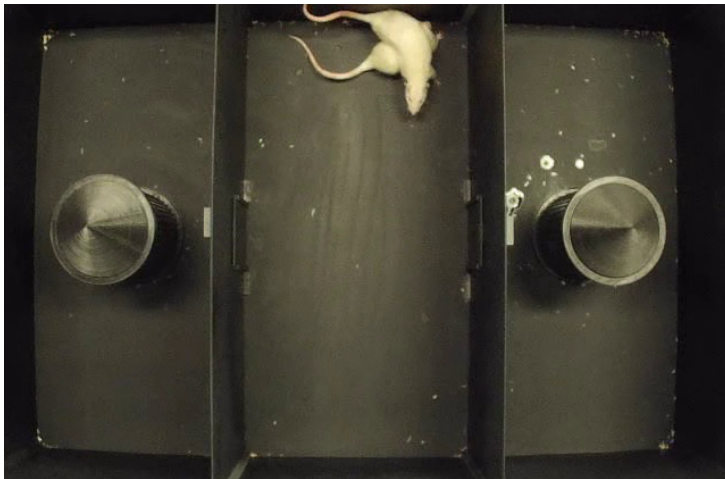


Figure 1. An example, where one of the rats is mounting the other one. The rats are highly similar and may heavily occlude each other from time to time giving rise to ambiguity in the instance segmentation and keypoint detection tasks due to hidden or heavily occluded regions.

2. Proposed methods

In this section, we first describe our proposed method for automatic annotation. Then we present the augmentation techniques and the architecture used for body part segmentation and keypoint detection.

The annotation process starts with the foreground-background segmentation (Section 2.1), then we annotate each foreground segment by combining several computer vision methods (Section 2.2). Figure 2 shows the pipeline of the system.

By augmenting the generated samples (Section 2.3) it is possible to train a deep model (Section 2.4) in a self-supervised manner.

2.1. Foreground-background segmentation

As the camera is stationary, it is possible to separate the foreground from the background by calculating the mode value of each pixel for multiple time steps then subtracting this value from the actual image.

Given a set of images $\mathcal{I} = \{x_1, x_2, \dots, x_N\}$ recorded by a stationary camera, where N is the number of images, and $x_i = [0, 255]^{H \times W \times 3}$ ($i = 1, \dots, N$) is a

colored (RGB) image. The background can be estimated by

$$B = \text{mode} \left(\begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_N \end{bmatrix}, \text{axis} = 0 \right).$$

Then the foreground of image i can be determined by

$$F_i = x_i - B \quad (i = 1, \dots, N).$$

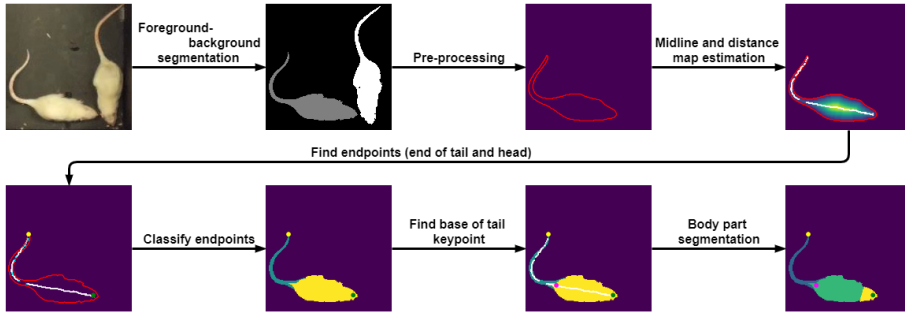


Figure 2. Pipeline of the automatic annotation method. The process starts with foreground-background segmentation; then, after pre-processing, medial axis transform is applied. From the endpoints of the midline, head and end of tail keypoints can be determined based on the area of their watershed segmentation. Then the base of the tail keypoint can be found by taking the median value of the distance transform. Finally, the body part segmentation is given by the watershed algorithm initiated from the acquired keypoints. Best viewed in color.

2.2. Automatic annotation

For automatically annotating the foreground segments with keypoints and body parts, we combined several computer vision methods. We distinguished three keypoints: (i) the head, (ii) the base of the tail, and (iii) the end of the tail; and three body parts: (i) the head, (ii) the body, and (iii) the tail. The pipeline of the overall algorithm consists of 6 steps (Figure 2):

1. Pre-processing: Improve initial segmentation and decide whether it is possible to annotate the mask.
2. Midline and distance map estimation: Estimate them by using medial axis transformation [17].

3. Find endpoints (end of tail and head): Find a pair of connected points via the longest minimum cost path given the midline.
4. Classify endpoints: By applying the watershed algorithm [12], the endpoints can be classified, given the area of the resulting segments.
5. Find base of tail keypoint: The last keypoint can be estimated by the location of the distance map's median value on the midline.
6. Body part segmentation: The body parts of the rat can be determined by using another watershed algorithm.

Pre-processing: The resulting masks of the foreground-background segmentation can be noisy. In order to use them, we need to do some pre-processing. To this end, we need to

1. Clean the mask: first we remove holes, then apply closing and finally remove torn off segments, then
2. Decide whether it is possible to annotate the mask: this can be checked by measuring its convexity, and finally
3. Smooth the boundary: otherwise, the resulting skeleton from the medial axis transform may contain spurious line segments and spoil our midline estimation. Therefore, we smooth the boundary by fitting 4th degree B-splines to it.

Midline and distance map estimation: After pre-processing, the skeleton and the distance transform of the mask are determined by the medial axis transform [17]. The distance transform is a grayscale image, where the intensity values represent the distance from the boundary, while the skeleton is a binary image, a remnant of the original mask. For each skeleton point, the corresponding distance transform value represents the radius of the largest circle, which is centered at that point and touches at least two boundary points. All other points of the circle are within or at the boundary of the mask. As this process is highly sensitive to irregularities in the boundary, proper smoothing is required to give us an accurate midline estimate.

We used the implementation of the medial axis transformation from the scikit-image Python package.

Find endpoints (end of tail and head): From the endpoints of the skeleton, we need to find those two, which are the furthest from each other. As the animals can take arbitrary poses, a simple pixel distance will not give optimal results. The endpoints should be selected by finding the longest minimum-cost path between them, given the midline.

Let $\mathcal{P} = \{p_1, p_2, \dots, p_K\}$ a set of all endpoints of the medial axis transform and $D = \mathbb{N}^{H \times W}$ the cost matrix. Each endpoint represents a coordinate of the image: $p_k = (h, w)$, where $h \in [1, H] \wedge w \in [1, W]$, and all endpoints have a cost of one:

$$D_{i,j} = \begin{cases} 1, & \text{if } (i, j) \in \mathcal{P} \\ \epsilon, & \text{otherwise} \end{cases},$$

where $1 < \epsilon \in \mathbb{N}$. In this case, the end of the tail and the head keypoints can be found using the following formula

$$(q, r) = \underset{p_i, p_j \in \mathcal{P} \wedge p_i \neq p_j}{\operatorname{argmax}} \operatorname{MCD}(p_i, p_j, D),$$

where $\operatorname{MCD}(., ., .)$ calculates the minimum-cost path between the specified coordinates given the cost matrix.

Classify endpoints: To determine which endpoint corresponds to the head and which one to the end of the tail, we apply the watershed [12] algorithm using the negative distance transform and initializing the basins with the two endpoints. Then we classify the point with the larger segmented area as the head and the other one as the end of the tail keypoint.

We used the implementation of the watershed algorithm from the scikit-image Python package.

Find base of tail keypoint: After we have acquired the head and the tail keypoints, the final keypoint location is determined by the position of the distance transform's median value on the midline that separates the larger and the smaller values.

We also perform an additional validation step in order to exclude conceivably incorrect annotations. By measuring the tail-to-body ratio [4], we can decide whether the tail is partially missing due to foreground-background segmentation mistakes or not and dismiss the mask if necessary.

Body part segmentation: Given all the keypoints, we can segment the body parts by applying another watershed algorithm.

Finally, we can do some post-processing steps to improve the results by placing the keypoints close to the nearest corner of the mask and trimming protruding parts of the head segment. Figure 3 shows the final result of the automatic annotation pipeline.

2.3. Augmentation

Since the proposed computer vision-based automatic annotation method (Section 2.2) is only capable of annotating rats if they are not occluded, we need

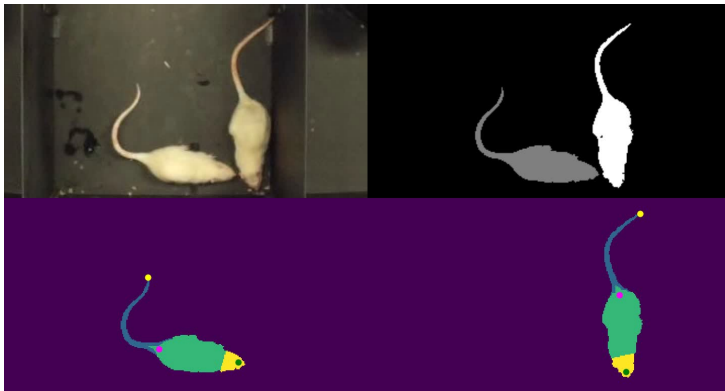


Figure 3. Result of the CV-based automatic annotation method. In the first row the input image and its foreground-background segmentation, and in the second row the annotated instances can be seen.

to introduce occlusions in the dataset by augmenting the generated training samples. Otherwise, the trained model could not exceed the performance of the initial method. To simulate occlusions, we randomly cut out one of the rats, fill its place in with the background, then shift the segment near the other instance. To achieve the best results, we combined several augmentation techniques commonly used in the video segmentation literature [8, 13]. We

- applied various kind of smoothing: Gaussian, median pyramid as well as no smoothing,
- rotated objects around its center by -45° to 45° ,
- scaled the shifted segment by 0.9 to 1.1, and
- deformed the mask via thin-plate-spline warping [2].

2.4. Architecture overview

Several deep architectures are capable of instance segmentation and key-point detection, such as [3, 16, 5]. For ease of use and modularity, we chose the Mask R-CNN architecture (see Figure 4). This model consists of 3 main parts:

- The backbone network is typically a ResNet50 [6] wrapped in a feature pyramid network (FPN) [10]. It accepts an RGB image as its input and extracts a set of feature maps with different spatial resolutions.
- The region proposal network (RPN) is inputted by the feature maps and generates bounding box proposals based on their objectness score, then

- The region of interest (ROI) heads process the proposed boxes and return the final results.

By combining different types of ROI heads, we can produce a model that is capable of instance segmentation and keypoint detection, and another one which is able to detect and segment body parts.

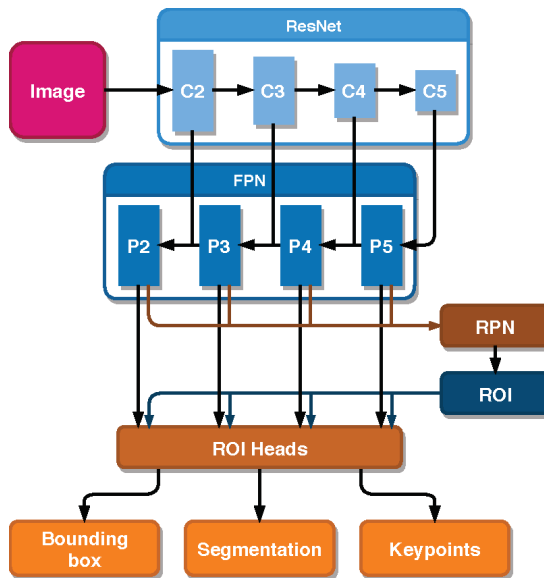


Figure 4. Mask R-CNN architecture. The network takes an RGB image, feeds it through the backbone, which extracts a feature pyramid, where each feature map has half the spatial resolution as the one before. Then the RPN makes bounding box proposals, which are processed by the ROI heads.

3. Results

Here, we introduce the dataset and the performance metric and present the results of our proposed methods.

3.1. Experimental setup

We implemented the automatic annotation method in Python using the scikit-image image processing library and NumPy. For the deep learning model, we used the Mask R-CNN implementation of Detectron2¹.

¹<https://github.com/facebookresearch/detectron2>

We used ResNet50 with FPN for the backbone, like in the original Mask R-CNN architecture [5]. We initialized the weights of our model with the COCO instance segmentation baseline in the case of body part segmentation, and with the COCO person keypoint detection baseline in the case of instance segmentation and keypoint detection. During training, we used a batch size of 512 for the heads and a batch size of 4 for the rest of the network. We initialized the learning rate to 0.00025 and set the object keypoint similarity (OKS) sigma values to 0.079, 0.107, and 0.089 for the head, base of the tail, and end of the tail keypoints, respectively. These sigma values correspond to the shoulders, hips, and ankles in the original COCO benchmark [9]. We trained each model for a maximum of 50,000 iterations on an NVIDIA GeForce GTX 1080 graphics card. During the evaluation, we set the confidence threshold to 70%. All other parameters were set to the default values.

Our implementation is available on GitHub².

3.2. Database

The database contains a continuous video of 2 white rats in a black container. The video was recorded by a stationary camera fixed on the top of the container. The video is 58994 frame-long, and the original resolution was 1280x720, from which we cropped a 640x420 region containing only the box. We only used the first 15,000 frames for training, and we annotated 200 frames, where no occlusion was present for the sake of evaluation.

The main challenges posed by the dataset are (i) the lack of ground truth labels, (ii) the highly similar objects, and (iii) the large amount of occlusion. (Figure 1.) To attain good results, one should address all of these problems.

Benchmark measure To measure the performance of our method, we use the metrics of the COCO Detection and Keypoint benchmark [9]. We evaluated the methods on the hand-annotated samples. We use average precision (AP) as our primary performance measure. AP is calculated by averaging over 10 Intersection over Union (IoU) levels from 0.5 to 0.95 with a step size of 0.05. The value of the AP is ranging between 0 and 1, and the higher, the better. It can also be calculated for keypoints, but in this case, the Object Keypoint Similarity (OKS) is used instead of IoU.

3.3. Automatic annotation

The proposed automatic annotation method is composed of several CV methods. Most of them can be tuned by setting their parameters or being replaced by alternative approaches, such as using active contours [7] instead of B-splines to smooth object boundary. We tried several combinations of parameters and methods from which a few are presented in Table 1.

²https://github.com/lkopi/rat_segmentation

We got the best overall result when we were using the method described in Section 2.2. On keypoint detection and instance segmentation, it reached 48.91 and 53.22 AP, respectively, while on body part segmentation, the AP was merely 9.38. Nevertheless, with proper augmentation, it proved to be enough for training deep neural networks. In the rest of the paper, we refer to this as baseline or CV-based approach. Figure 3 shows the result of this method.

| Set head to nearest corner | Clean mask | Trim protruding parts | Bounding box AP | Keypoint AP | Segmentation AP | Bounding box AP | Segmentation AP |
|----------------------------------|---------------|-----------------------------|--------------------|----------------|--------------------|--------------------|--------------------|
| ✓ | ✗ | ✓ | 56.76 | 47.36 | 44.00 | 11.81 | 9.37 |
| ✗ | ✗ | ✓ | 68.13 | 47.76 | 53.03 | 16.59 | 9.37 |
| ✗ | ✓ | ✓ | 67.71 | 48.91 | 53.22 | 17.40 | 9.38 |
| ✗ | ✓ | ✗ | 67.90 | 48.83 | 53.42 | 16.20 | 9.38 |

Table 1. Results of the automatic annotation method. The first three columns show some of the system’s parameters. The subsequent three and the last two columns show the performance on the keypoint detection and instance segmentation task, and on the body part segmentation task. We achieved the best overall performance with the method highlighted in bold. The details are described in Section 2.2.

3.4. Self-supervised method

We trained two separate models on the data produced by the automatic annotation method. The task of the first one was keypoint detection and instance segmentation, while the second model was trained to detect and segment body parts. The hand-annotated labels were not introduced in the training process; they were only used to evaluate the performance of the final models. On both tasks, the trained models outperformed the CV-based approach. On keypoint detection, it was better by 35.25%, while on body part segmentation, it surpassed the baseline by 2.35 times, contributing to a 135% improvement.

Figure 5 shows some outputs of the final models. Both models perform well even in annotating heavily occluded objects, but the keypoint detection model cannot detect objects when they are very close to each other. This issue is due to the architecture because Mask R-CNN uses non-maximum suppression to choose the best fitting bounding box for each object. Nevertheless, in such cases, the body part model can manage the segmentation of the object.

3.4.1. Keypoint detection and instance segmentation model

After 30K iterations, the performance of the networks, which used pyramid median or Gaussian smoothing, stagnated. While without smoothing, we were able to train the model for 50K iterations. In all cases, we observed that by applying more augmentation, better results could be reached. We got the best outcomes when no smoothing was applied. However, these models needed much more iterations to compete with others. By trained only for 30K iter-

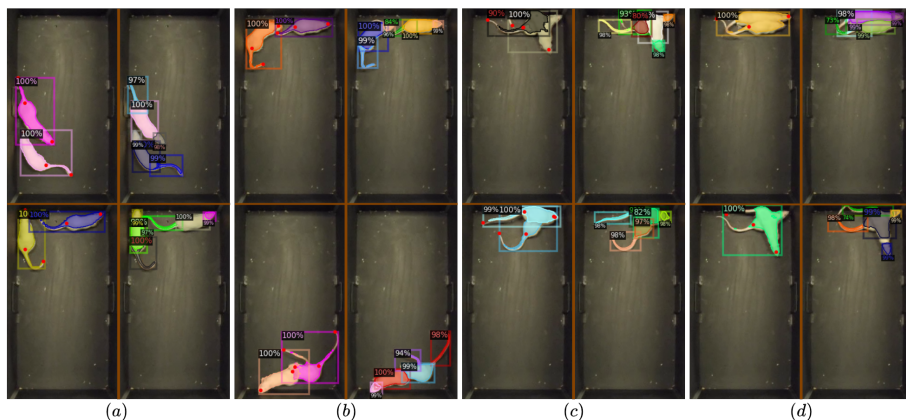


Figure 5. Result of the final method. The left side of each image shows the result of the keypoint detection and instance segmentation model, and the right side shows the result of the body part segmentation model. The columns address the different amount of occlusion present on the scene. If there is no occlusion (a) or just partial occlusion (b) is present on the scene, both methods perform well. If the objects are heavily occluded (c), the quality of the keypoint model starts degrading. After a point, (d), it cannot separate instances due to limitations of the Mask R-CNN architecture, but the body part model can still segment the objects.

ations, their performance was worse than using just basic augmentation. It reached only 82.57, 47.79, and 60.09 AP. However, after training for 20K more iterations, it outperformed all previous models. The final results are 90.27, 61.92, and 77.53 AP on object detection, keypoint detection, and instance segmentation, respectively. Compared to the baseline (CV method), it achieved a 35.25% improvement. We summarized the results in Table 2.

3.4.2. Body part segmentation model

In the case of body part segmentation (Table 3), we observed the clues of over-fitting in a much earlier state. Models trained only for 5K iterations outperformed all subsequent stages. We discovered similar insights as in the keypoint detection case. In general, models trained with tedious augmentation achieved similar or slightly better results except in the no smoothing experiments, where the performance decreased when thin-plate-spline warping was applied. The added value of augmentation was modest ($\pm(0-6)\%$) as the applied techniques do not affect parts significantly. Consequently, more augmentation might be needed to reach better results.

We reached the best performance when only rotation and scaling were used

| Rotation and scaling | Warping | Smoothing | # of iterations | Bounding box AP | Keypoint AP | Segmentation AP |
|----------------------|--------------|-----------|-----------------|-----------------|--------------|-----------------|
| \times | \times | None | 30K | 84.04 | 58.75 | 60.77 |
| \checkmark | \times | None | 50K | 85.30 | 61.30 | 63.46 |
| \checkmark | \checkmark | None | 50K | 90.27 | 61.92 | 77.53 |
| \times | \times | Pyramid | 30K | 83.98 | 53.01 | 60.56 |
| \checkmark | \times | Pyramid | 30K | 85.58 | 51.47 | 59.01 |
| \checkmark | \checkmark | Pyramid | 30K | 83.43 | 56.13 | 59.45 |
| \times | \times | Gauss | 30K | 83.17 | 43.13 | 59.96 |
| \checkmark | \times | Gauss | 30K | 84.26 | 55.99 | 60.12 |
| \checkmark | \checkmark | Gauss | 30K | 85.30 | 61.73 | 60.58 |

Table 2. Results of the trained model on the keypoint detection and instance segmentation task. We studied the effect of different augmentation techniques presented in Section 2.3. Warping denotes thin-plate-spline warping with a keypoint shift of a maximum of 10 pixels. The best performance was acquired when no smoothing was applied. In all cases, the more augmentation we used, the better the result got.

without any smoothing. The model achieved 34.01 AP on body part detection and 28.87 AP on segmentation, which accounts for a 134.78% gain compared to the CV-based approach.

| Rotation and scaling | Warping | Smoothing | # of iterations | Bounding box AP | Segmentation AP |
|----------------------|--------------|-----------|-----------------|-----------------|-----------------|
| \times | \times | None | 5K | 30.50 | 28.77 |
| \checkmark | \times | None | 5K | 34.01 | 28.87 |
| \checkmark | \checkmark | None | 5K | 32.27 | 29.17 |
| \times | \times | Pyramid | 5K | 32.60 | 28.98 |
| \checkmark | \times | Pyramid | 5K | 32.32 | 27.87 |
| \checkmark | \checkmark | Pyramid | 5K | 33.07 | 29.10 |
| \times | \times | Gauss | 5K | 32.53 | 27.93 |
| \checkmark | \times | Gauss | 5K | 32.23 | 27.88 |
| \checkmark | \checkmark | Gauss | 5K | 33.13 | 27.51 |

Table 3. Results on the body part segmentation tasks. We applied similar augmentation techniques to the keypoint detection task. As the results show, their contribution was smaller as they did not significantly affect individual components. During training, the performance of the models dwindled after 5,000 iterations, so more rigorous augmentation might be necessary to further improve the AP. The best overall result was achieved when only rotation and scaling was used. This method is marked in bold typefaces.

4. Summary

In this paper, we presented a method capable of instance segmentation, keypoint detection, and body part segmentation without the need for any hand-annotated data. We studied the effect of different augmentation techniques and achieved AP of 61.92%, 77.53%, and 28.87% on instance segmentation, keypoint detection, and body part segmentation, respectively.

Our results show that automatic annotation of rat images is possible without the need for manual labeling. The method presented here can be used to analyze the interaction between instances and to speed up the annotation process if more precise labels are needed. To improve the performance, one should address the shortcoming of Mask R-CNN in the case of instance segmentation, which is due to non-maximum suppression, by changing the architecture [3] or by separating bounding boxes if more than one instance is present.

The system could be extended to handle video sequences via bipartite matching [1] or by incorporating optical flow estimated by deep neural networks [14], among others. In addition, the power of such temporal extension methods and the fusion of the keypoint and the body part models may be able to track such very similar and flexible objects subject to heavy occlusions in order to study their interactions in an automated fashion.

Author contributions. László Kopácsi and András Lőrincz conceived and designed the research, László Kopácsi and Áron Fóthi performed computational analyses.

This work is a detailed version of a MaCS 2020 presentation.

References

- [1] **Bergmann, P., T. Meinhardt, and L. Leal-Taixe**, Tracking without bells and whistles, in: *Proceedings of the IEEE International Conference on Computer Vision*, (2019) 941–951.
- [2] **Bookstein, F.L.**, Principal warps: Thin-plate splines and the decomposition of deformations, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(6) (1989), 567–585.
- [3] **Carion, N., F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko**, End-to-end object detection with transformers, arXiv preprint (2020) <https://arxiv.org/abs/2005.12872>
- [4] **Hanson, F.B., and F. Heys**, Correlations of body weight, body length, and tail length in normal and alcoholic albino rats, *Genetics*, **9**(4) (1924), 368.
- [5] **He, K., G. Gkioxari, P. Dollár, and R. Girshick**, Mask r-cnn, in: *Proceedings of the IEEE International Conference on Computer Vision* (2017), 2961–2969.
- [6] **He, K., X. Zhang, S. Ren, and J. Sun**, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), 770–778.
- [7] **Kass, M., A. Witkin, and D. Terzopoulos**, Snakes: Active contour models, *International Journal of Computer Vision*, **1**(4) (1988), 321–331.

- [8] **Khoreva, A., R. Benenson, E. Ilg, T. Brox, and B. Schiele**, Lucid data dreaming for video object segmentation, *International Journal of Computer Vision*, (2018), 1–23.
- [9] **Lin, T.-Y., M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C.L. Zitnick**, Microsoft coco: Common objects in context, in: *European Conference on Computer Vision*, (2014), 740–755.
- [10] **Lin, T.-Y., P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie**, Feature pyramid networks for object detection, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2017), 2117–2125.
- [11] **Mathis, A., P. Mamidanna, K.M. Cury, T. Abe, V.N. Murthy, M.W. Mathis, and M. Bethge**, DeepLabCut: markerless pose estimation of user-defined body parts with deep learning, *Nature Neuroscience*, **21** (2018), 1281–1289.
- [12] **Neubert, P., and P. Protzel**, Compact watershed and preemptive slic: On improving trade-offs of superpixel segmentation algorithms, in: *Proceedings of the 22nd International Conference on Pattern Recognition, IEEE*, (2014), 996–1001.
- [13] **Perazzi, F., A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung**, Learning video object segmentation from static images, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2017), 2663–2672.
- [14] **Sun, D., X. Yang, M.-Y. Liu, and J. Kautz**, Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2018), 8934–8943.
- [15] **Varga, V., and A. Lőrincz**, Reducing human efforts in video segmentation annotation with reinforcement learning, *Neurocomputing*, **405** (2020), 247–258.
- [16] **Wang, J., K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu et al.**, Deep high-resolution representation learning for visual recognition, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (2020).
- [17] **Zhang, T.Y., and C.Y. Suen**, A fast parallel algorithm for thinning digital patterns, *Communications of the ACM*, **27(3)** (1984), 236–239.

Á. Fóthi, L. Kopácsi and A. Lőrincz

Department of Artificial Intelligence

Faculty of Informatics

Eötvös Loránd University

Budapest

Hungary

fa2@inf.elte.hu, kopacsi@inf.elte.hu and lorincz@inf.elte.hu