



Heatmap Regression for Lesion Detection Using Pointwise Annotations

Chelsea Myers-Colet¹(✉), Julien Schroeter¹, Douglas L. Arnold²,
and Tal Arbel¹

¹ Centre for Intelligent Machines, McGill University, Montreal, Canada
{cmyers,julien,arbel}@cim.mcgill.ca

² Montreal Neurological Institute, McGill University, Montreal, Canada
douglas.arnold@mcgill.ca

Abstract. In many clinical contexts, detecting all lesions is imperative for evaluating disease activity. Standard approaches pose lesion detection as a segmentation problem despite the time-consuming nature of acquiring segmentation labels. In this paper, we present a lesion detection method which relies only on point labels. Our model, which is trained via heatmap regression, can detect a variable number of lesions in a probabilistic manner. In fact, our proposed post-processing method offers a reliable way of directly estimating the lesion existence uncertainty. Experimental results on Gad lesion detection show our point-based method performs competitively compared to training on expensive segmentation labels. Finally, our detection model provides a suitable pre-training for segmentation. When fine-tuning on only 17 segmentation samples, we achieve comparable performance to training with the full dataset.

Keywords: Lesion detection · Lesion segmentation · Heatmap regression · Uncertainty · Multiple sclerosis

1 Introduction

For many diseases, detecting the presence and location of all lesions is vital for estimating disease burden and treatment efficacy. In stroke patients, for example, the location of a cerebral hemorrhage was shown to be an important factor in assessing the risk of aspiration [1] thus, failing to locate even a single one could drastically impact the assessment. Similarly, in patients with Multiple Sclerosis (MS), detecting and tracking all gadolinium-enhancing lesions (Gad lesions), whether large or small, is especially relevant for determining treatment response in clinical trials [2]. Detecting all Gad lesions is imperative as just one new lesion indicates new disease activity.

To achieve this goal, standard practice in deep learning consists of training a lesion segmentation model with a post-processing detection step [3, 4]. However, segmentation labels are expensive and time consuming to acquire. To this end, we

develop a lesion detection model trained on pointwise labels thereby reducing the manual annotation burden. Unlike previous point annotation-based methods [5–7], ours combines the ability to detect a variable number of lesions with the benefit of leveraging a probabilistic approach. Indeed, our refinement method is not only independent of a specific binarization threshold, it offers a unique way of estimating the lesion existence probability. Our contributions are threefold:

- (1) We demonstrate the merit of training on point annotations via heatmap regression over segmentation labels on the task of Gad lesion detection. With weaker labels, our models still achieve better detection performance.
- (2) Our proposed refinement method allows for a reliable estimation of lesion existence uncertainty thus providing valuable feedback for clinical review.
- (3) When the end goal is segmentation, our detection models provide a suitable pre-training for fine-tuning on a limited set of segmentation labels. When having access to only 17 segmentation samples, we can achieve comparable performance to a model trained on the entire segmentation dataset.

2 Related Work

Point annotations are often extremely sparse which leads to instability during training of deep neural networks. Therefore, most state-of-the-art methods rely on the application of a smoothing operation to point labels. A Gaussian filter is commonly applied to create a heatmap as was done in [5, 6] for suture detection. Others have found success applying distance map transformations. For instance, Han et al. [8] and van Wijnen et al. [9] used Euclidean and Geodesic distance maps to perform lesion detection. We demonstrate the benefits of training with Gaussian heatmaps over distance maps as they offer a more precise and interpretable probabilistic prediction yielding superior detection performance.

Irrespective of the choice of smoothing used for training, detection methods will often differ in their post-processing refinement step, i.e. in extracting lesion coordinates from a predicted heatmap. The simplest approach consists in finding the location with the maximum mass [5, 10, 11] or computing the centre of mass [6]. Although these approaches easily allow for the detection of multiple lesions, they require careful tuning of the binarization threshold and are susceptible to missing both isolated and overlapping peaks. More sophisticated methods exist which aim to fit a Gaussian distribution to the predicted heatmap thus retaining its probabilistic interpretation, e.g. [12, 13]. Specifically, to perform cephalogram landmark detection Thaler et al. [7] align a Gaussian distribution via Least Squares curve fitting. Since the approach taken in [7] is limited to a set number of landmarks, we extend it to detect a variable number of lesions. Our method thus offers the flexibility of simpler approaches, without any dependence on a binarization threshold, while providing a probabilistic interpretation.

3 Method

In this work, we propose a strategy to detect the presence and location of multiple lesions from brain MRIs of patients with a neurodegenerative disease. Our model




Fig. 1. Overview of detection method given a predicted heatmap $\hat{\mathbf{H}}_i$: lesion candidates are found by (1) locating the global maximum, (2) fitting a Gaussian distribution to an extracted region and (3) subtracting the influence of this lesion from the heatmap. (4) Repeat steps 1–3 before (5) filtering out unlikely lesions.

is trained via heatmap regression (Sect. 3.1) while lesion detection is performed in a post-processing step (Sect. 3.2). Finally, we present a transfer learning scheme to perform segmentation on a limited dataset (Sect. 3.3).

3.1 Training via Heatmap Regression

The proposed heatmap regression training scheme requires a domain expert to label only a single point identifying each lesion, e.g. by marking the approximate centre of the lesion. To stabilize training, we apply a Gaussian filter with smoothing parameter σ to the point annotations thus creating a multi-instance heatmap [5–7, 14–16]. Since all lesions are represented by a single point and smoothed using the same value of σ , equal importance is attributed to lesions of all sizes. We train a model f_θ to map a sequence of input MRIs to a predicted heatmap $\hat{\mathbf{H}}_i$.

3.2 Detection During Inference

Given a continuous heatmap $\hat{\mathbf{H}}_i$, we now aim to detect individual lesions. Specifically, for patient i , we wish to represent the k^{th} detected lesion by a single point, $\hat{\mu}_{ik}$, which can be extracted from the heatmap. We assume the predicted heatmap $\hat{\mathbf{H}}_i$ will model a sum of Gaussian distributions (each describing a single lesion) to mimic the target heatmap:¹

$$\hat{\mathbf{H}}_i = \sum_K \hat{\mathbf{H}}_{ik} = \sum_K \mathcal{N}(\hat{\mu}_{ik}, \sigma) \quad (1)$$

¹ Valid as long as f_θ sufficiently minimizes the loss and thus models the target.

Our method essentially aims to find the individual Gaussian distributions comprising the sum in Eq. 1 in an iterative manner as shown in Fig. 1. We now describe each depicted step in detail.

(1) Locate Global Maximum. The location of the global maximum serves as an initial estimate for the k^{th} predicted lesion centre, $\hat{\boldsymbol{\mu}}_{ik}$.

(2) Gaussian Fitting. In the region surrounding a detected lesion with centre $\hat{\boldsymbol{\mu}}_{ik}$, we fit a Gaussian distribution with normalizing constant $\hat{\alpha}_{ik}$:

$$\hat{\mathbf{H}}_{ik} = \hat{\alpha}_{ik} \mathcal{N}(\hat{\boldsymbol{\mu}}_{ik}, \sigma) \quad (2)$$

Provided there is minimal overlap between neighbouring lesions, we can use a Least Squares curve fitting algorithm to estimate $\hat{\alpha}_{ik}$ and $\hat{\boldsymbol{\mu}}_{ik}$.² The normalizing constant, $\hat{\alpha}_{ik}$, represents the prior probability of producing a peak in this region (from Bayes' Theorem), i.e. it is the belief that a lesion exists in the given region. We thus refer to $\hat{\alpha}_{ik}$ as the lesion existence probability (similar to [17]). As an initial estimate for $\hat{\alpha}_{ik}$, we sum within the extracted region, i.e. the hypothesis space, as shown in Fig. 1 (2).

(3) Subtract. Now that potential lesion k has been identified and fitted with a continuous Gaussian function, we remove its contribution to the sum in Eq. 1. This allows our method to more easily detect the individual contributions of neighbouring lesions with overlapping Gaussian distributions.

$$\hat{\mathbf{H}}'_i = \hat{\mathbf{H}}_i - \hat{\alpha}_{ik} \mathcal{N}(\hat{\boldsymbol{\mu}}_{ik}, \sigma) \quad (3)$$

(4) Repeat. Since we have subtracted the contribution of lesion k from the aggregated heatmap, the global maximum now corresponds to a different candidate lesion. Steps 1 to 3 are repeated until a maximum number of lesions have been found or when the lesion existence probability drops below a threshold, e.g. 0.01.

(5) Filtering. Lesions with a low probability of existence are discarded (threshold optimized on the validation set). By overestimating the lesion count and subsequently discarding regions unlikely to contain a lesion, we can better account for noisy peaks in the heatmap. We evaluate the calibration of these probabilities and demonstrate the validity of this filtering step (Sect. 4.2).

3.3 Segmentation Transfer Learning

In addition to detection, estimating a lesion segmentation can be beneficial for assessing lesion load. We therefore design a transfer learning scheme which first relies on building a strong lesion detector using point annotations before fine-tuning on a small segmentation dataset. Specifically, we (1) train a detection model on point annotations until convergence; (2) build a small segmentation training set; (3) fine-tune the detection model on segmentation samples only. Training in this manner minimizes the amount of detailed segmentation labels that must be generated.

² Valid for Gad lesions given a sufficiently small smoothing parameter σ .

4 Experiments and Results

The proposed heatmap regression model is compared against three benchmarks in terms of detection performance. We train models on (1) segmentation labels, (2) Euclidean distance maps and (3) Geodesic distance maps [8,9]. Similar to our method, lesions are detected from the output prediction in a post-processing step. Here, we instead binarize the output at threshold τ (optimized on the validation set), cluster connected components to form detected lesions and use the centre of mass (segmentation) or the maximum (detection) to represent the lesion (referred to as CC). As an additional benchmark, we apply this method to heatmap outputs from our proposed regression models. This is in line with detection methods used by [18,19] for segmentation outputs and [5,6] for heatmap predictions.

4.1 Experimental Setup

Dataset. We evaluate our method on Gad lesion detection as they are a relevant indicator of disease activity in MS patients [20]. However, their subtlety and extreme size variation makes them difficult to identify. Experiments are performed using a large, multi-centre, multi-scanner proprietary dataset consisting of 1067 patients involved in a clinical trial to treat Relapsing-Remitting MS. Multi-modal MRIs, including post-contrast T1-weighted MRI, are available for each patient and are provided as inputs to our system. For fairness, we create train (60%), validation (20%) and test (20%) sets by first splitting at the patient level. We have access to manually derived Gad lesion segmentation masks. Each sample is first independently rated by two experts who then meet to produce a consensus. Point labels were generated directly from segmentation masks by calculating the centre of mass of each lesion and transformed into either heatmaps, using a Gaussian kernel with smoothing parameter σ , or distance maps (baseline methods), using decay parameter p . Hyperparameters were selected based on validation performance.

Model. We train a modified 5-layer U-Net [21] with dropout and instance normalization using a Mean-Squared Error loss for heatmap regression and a weighted cross-entropy loss for segmentation. See code for details³.

Evaluation. We apply the Hungarian algorithm [22] to match predicted lesions to ground truth lesions using Euclidean distance as a cost metric. Assignments with large distances are considered both a false positive and a false negative.

³ <https://github.com/ChelseaM-C/MICCAI2022-Heatmap-Lesion-Detection>.

Table 1. Lesion detection results as a mean over 3 runs. Reported is the detection F1-score, precision, recall and small lesion recall for models trained with segmentation, Gaussian heatmap or distance map (Geodesic, Euclidean) [8,9] labels using connected components (CC) or Gaussian fitting (GAUSSIAN).

Label type	Detection method	F1-score	Precision	Recall	Small lesion recall
Segmentation	CC	85.4 ± 0.02	85.3 ± 1.10	85.5 ± 1.15	67.7 ± 3.58
Euclidean map	CC	80.6 ± 1.01	92.6 ± 1.28	71.4 ± 2.14	51.0 ± 3.28
Geodesic map	CC	73.7 ± 4.89	81.0 ± 7.97	67.8 ± 2.98	47.8 ± 2.81
Gaussian heatmap	CC	83.9 ± 0.27	80.9 ± 4.43	87.3 ± 2.64	75.0 ± 5.75
	GAUSSIAN	86.3 ± 0.24	87.0 ± 1.89	85.7 ± 1.44	70.4 ± 4.47

4.2 Lesion Detection Results

Despite only having access to point annotations, the proposed Gaussian heatmap approach performs competitively with the segmentation baseline (see Table 1). In fact, our proposed iterative detection method (GAUSSIAN) even slightly outperforms the segmentation model on all detection metrics. By contrast, both distance map approaches show notably worse performance with especially low recall scores indicating a high number of missed lesions. The proposed model additionally outperforms competing methods for the task of small lesion detection (3 to 10 voxels in size) underlining the merit of training directly for detection. Segmentation models will typically place more importance on larger lesions since they have a higher contribution to the loss, a bias not imposed by our detection model. Our model additionally does not sacrifice precision for high recall on small lesions; we perform on par with segmentation.

Lesion Existence Probability Evaluation. We evaluate the quality of our fitted lesion existence probabilities on the basis of calibration and derived uncertainty to justify both the curve fitting and filtering steps.

(1) *Calibration.* We compare the calibration [23] of the lesion existence probabilities before and after Least Squares curve fitting. Recall the initial estimate for α_{ik} is found by summing locally within the extracted region. Our proposed existence probabilities are well calibrated (Fig. 2a), with little deviation from the ideal case thus justifying our proposed filtering step. As well, the fitted probabilities are significantly better calibrated than the initial estimates demonstrating the benefit of curve fitting.

(2) *Uncertainty.* While it is important to produce accurate predictions, quantifying their uncertainty is of equal importance in the medical domain. We can compute the entropy of our lesion existence probabilities without sampling and show it is well correlated with detection accuracy. As we consider only the least uncertain instances, we observe a monotonic increasing trend, even achieving an accuracy of 100% (Fig. 2b). Similar results are achieved with the more standard MC Dropout approach [24] applied to segmentation outputs (calculated at a lesion level as in [18]).




Fig. 2. Lesion existence probability evaluation. (a) Calibration of unfitted (pink) vs. fitted (green) probabilities. (b) Detection accuracy of least uncertain samples according to our model (green) vs. MC Dropout applied to segmentation (pink). (Color figure online)

Our derived lesion existence probabilities are not only well-calibrated, they produce meaningful uncertainty estimates. With only a single forward pass, our uncertainty estimates perform on par with standard sampling-based approaches.

4.3 Lesion Segmentation via Transfer Learning

To demonstrate the adaptability of our method, we fine-tune the trained heatmap regression models with a small segmentation dataset as described in Sect. 3.3. Specifically, we use segmentation labels for a randomly chosen 1% of our total training set for fine-tuning. To account for bias in the selected subset, we repeat this process 3 times and average the results. For comparison, we train from scratch with this limited set as well as on the full segmentation training set. We additionally include results on random subsets of 5% and 10% in the appendix along with the associated standard deviation of each experiment.

Remarkably, our pre-trained models show only a 3% drop in segmentation F1-score performance with a mere 1% of the segmentation labels compared to the model trained on the full segmentation dataset (see Table 2). By contrast, the model trained from scratch with the same 1% of segmentation labels shows a 10% drop in segmentation F1-score. This emphasizes the importance of detecting lesions since models trained from scratch in the low data regime show considerably lower detection F1-score. It is clear the models do not require very much data in order to properly segment lesions as demonstrated by competitive performance of our pre-trained models. However, as indicated by poor detection performance of the pure segmentation models in the low data regime, it is clear these models need help localizing lesions before they can be segmented. We can

Table 2. Segmentation transfer learning results averaged over 3 random subsets. We present segmentation F1-score (SEG F1.) and detection (DET.) metrics on the fine-tuned segmentation models: F1-score, precision, recall. Pre-trained models are distinguished by their smoothing hyperparameter σ .

Quantity seg. labels	Pre-trained model	SEG. F1	DET. F1	DET. precision	DET. recall
100%	None	70.5 \pm 0.31	85.4 \pm 0.02	85.3 \pm 1.10	85.5 \pm 1.15
1%	None	60.7 \pm 2.48	69.7 \pm 5.92	79.0 \pm 6.47	62.7 \pm 7.69
	$\sigma = 1.0$	67.2 \pm 0.71	85.4 \pm 0.62	83.9 \pm 1.74	87.0 \pm 2.40
	$\sigma = 1.25$	67.6 \pm 1.31	84.5 \pm 0.38	86.0 \pm 0.67	83.1 \pm 1.32
	$\sigma = 1.5$	67.2 \pm 1.51	85.0 \pm 0.50	83.6 \pm 1.01	86.6 \pm 1.10
	Euclidean	66.3 \pm 0.18	84.7 \pm 1.45	83.3 \pm 3.24	86.1 \pm 1.67
	Geodesic	59.9 \pm 1.44	77.9 \pm 4.36	85.6 \pm 7.33	71.7 \pm 3.49

see a similar trend with the models pre-trained on distance maps. The Euclidean distance maps offered higher detection scores than the Geodesic ones (Table 1) and therefore serve as a better pre-training for segmentation, although still lower than our models.

5 Discussion and Conclusion

In this work, we have demonstrated how training a heatmap regression model to detect lesions can achieve the same, and at times better, detection performance compared to a segmentation model. By requiring clinicians to indicate a single point within each lesion, our approach significantly reduces the annotation burden imposed by deep learning segmentation methods. Our proposed method of iteratively fitting Gaussian distributions to a predicted heatmap produces well-calibrated existence probabilities which capture the underlying uncertainty.

Perhaps most significantly, our transfer learning experiments have revealed an important aspect about segmentation models. Our results demonstrate that segmentation models must learn first and foremost to find lesions. Indeed, our models, which are already adept at lesion detection, can easily learn to delineate borders with only a few segmentation samples. By contrast, the models provided with the same limited set of segmentation labels trained from scratch fail primarily to detect lesions thus lowering their segmentation scores. It therefore presents an unnecessary burden on clinicians to require them to manually segment large datasets in order to build an accurate deep learning segmentation model.

Although we have demonstrated many benefits, Gaussian heatmap matching has its limitations. The smoothing hyperparameter requires careful tuning to both maintain stable training and to avoid a significant overlap in peaks (especially for densely packed lesions). As well, the method still requires an expert annotator to mark approximate lesion centres however, this is much less time-consuming than fully outlining each lesion. We recognize this could introduce high variability in the labels regarding where the point is placed within each lesion. Though the current model was trained on precise centres of mass, the

proposed method does not necessarily impose any such constraints, in theory. Future work is needed to evaluate the robustness of the model to high variability in the label space.

In summary, our proposed training scheme and iterative Gaussian fitting post-processing step constitute an accurate and label-efficient method of performing lesion detection and segmentation.

Acknowledgements. This work was supported by awards from the International Progressive MS Alliance (PA-1412-02420), the Canada Institute for Advanced Research (CIFAR) Artificial Intelligence Chairs program (Arbel), the Canadian Natural Science and Engineering Research Council (CGSM-NSERC-2021-Myers-Colet) and the Fonds de recherche du Québec (303237). The authors would also like to thank Justin Szeto, Kirill Vasilevski, Brennan Nichyporuk and Eric Zimmermann as well as the companies who provided the clinical trial data: Biogen, BiOMS, MedDay, Novartis, Roche/Genentech, and Teva. Supplementary computational resources were provided by Calcul Québec, WestGrid, and Compute Canada.

References

1. Daniels, S.K., Foundas, A.L.: Lesion localization in acute stroke. *J. Neuroimaging* **9**(2), 91–98 (1999)
2. Rudick, R.A., Lee, J.-C., Simon, J., Ransohoff, R.M., Fisher, E.: Defining interferon β response status in multiple sclerosis patients. *Ann. Neurol. Official J. Am. Neurol. Assoc. Child Neurol. Soc.* **56**(4), 548–555 (2004)
3. Lundervold, A.S., Lundervold, A.: An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik* **29**(2), 102–127 (2019)
4. Doyle, A., Elliott, C., Karimghaloo, Z., Subbanna, N., Arnold, D.L., Arbel, T.: Lesion detection, segmentation and prediction in multiple sclerosis clinical trials. In: Crimi, A., Bakas, S., Kuijff, H., Menze, B., Reyes, M. (eds.) *BrainLes 2017*. LNCS, vol. 10670, pp. 15–28. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75238-9_2
5. Sharan, L., et al.: Point detection through multi-instance deep heatmap regression for sutures in endoscopy. *Int. J. Comput. Assist. Radiol. Surg.* **16**(12), 2107–2117 (2021). <https://doi.org/10.1007/s11548-021-02523-w>
6. Stern, A., et al.: Heatmap-based 2d landmark detection with a varying number of landmarks. arXiv preprint [arXiv:2101.02737](https://arxiv.org/abs/2101.02737) (2021)
7. Thaler, F., Payer, C., Urschler, M., Stern, D.: Modeling annotation uncertainty with gaussian heatmaps in landmark localization. arXiv preprint [arXiv:2109.09533](https://arxiv.org/abs/2109.09533) (2021)
8. Han, X., Zhai, Y., Yu, Z., Peng, T., Zhang, X.-Y.: Detecting extremely small lesions in mouse brain MRI with point annotations via multi-task learning. In: Lian, C., Cao, X., Rekić, I., Xu, X., Yan, P. (eds.) *MLMI 2021*. LNCS, vol. 12966, pp. 498–506. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87589-3_51
9. van Wijnen, K.M.H., et al.: Automated lesion detection by regressing intensity-based distance with a neural network. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11767, pp. 234–242. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32251-9_26
10. Donné, S., De Vylder, J., Goossens, B., Philips, W.: Mate: machine learning for adaptive calibration template detection. *Sensors* **16**(11), 1858 (2016)

11. Chen, B., Xiong, C., Zhang, Q.: CCDN: checkerboard corner detection network for robust camera calibration. In: Chen, Z., Mendes, A., Yan, Y., Chen, S. (eds.) ICIRA 2018. LNCS (LNAI), vol. 10985, pp. 324–334. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-97589-4_27
12. Zhang, F., Zhu, X., Dai, H., Ye, M., Zhu, C.: Distribution-aware coordinate representation for human pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7093–7102 (2020)
13. Graving, J.M., et al.: Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. *Elife* **8**, e47994 (2019)
14. Wang, X., Bo, L., Fuxin, L.: Adaptive wing loss for robust face alignment via heatmap regression. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 6971–6981 (2019)
15. Pfister, T., Charles, J., Zisserman, A.: Flowing convnets for human pose estimation in videos. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1913–1921 (2015)
16. Hervella, Á.S., Rouco, J., Novo, J., Penedo, M.G., Ortega, M.: Deep multi-instance heatmap regression for the detection of retinal vessel crossings and bifurcations in eye fundus images. *Comput. Methods Programs Biomed.* **186**, 105201 (2020)
17. Schroeter, J., Myers-Colet, C., Arnold, D., Arbel, T.: Segmentation-consistent probabilistic lesion counting. *Med. Imaging Deep Learn.* (2022)
18. Nair, T., Precup, D., Arnold, D.L., Arbel, T.: Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation. *Med. Image Anal.* **59**, 101557 (2020)
19. De Moor, T., Rodriguez-Ruiz, A., Mérida, A.G., Mann, R., Teuwen, J.: Automated lesion detection and segmentation in digital mammography using a u-net deep learning network. In: 14th International Workshop on Breast Imaging (IWBI 2018), vol. 10718, p. 1071805. International Society for Optics and Photonics (2018)
20. McFarland, H.F., et al.: Using gadolinium-enhanced magnetic resonance imaging lesions to monitor disease activity in multiple sclerosis. *Ann. Neurol.* **32**(6), 758–766 (1992)
21. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
22. Kuhn, H.W.: The Hungarian method for the assignment problem. *Naval Res. Logist. Q.* **2**(1–2), 83–97 (1955)
23. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International Conference on Machine Learning, pp. 1321–1330. PMLR (2017)
24. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: International Conference on Machine Learning, pp. 1050–1059. PMLR (2016)