



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Progressive Growing of Patch Size: Resource-Efficient Curriculum Learning for Dense Prediction Tasks

Stefan M. Fischer<sup>1,2,3</sup>, Lina Felsner<sup>1,2</sup>, Richard Osuala<sup>1,2,4</sup>, Johannes Kiechle<sup>1,2,3</sup>, Daniel M. Lang<sup>1,2</sup>, Jan C. Peeken<sup>1,2</sup>, and Julia A. Schnabel<sup>1,2,3,5</sup>

<sup>1</sup> Technical University Munich, Germany

<sup>2</sup> Helmholtz Munich, Germany

<sup>3</sup> Munich Center of Machine Learning (MCML), Germany

<sup>4</sup> Universitat de Barcelona, Spain

<sup>5</sup> King's College London, UK

**Abstract.** In this work, we introduce Progressive Growing of Patch Size, a resource-efficient implicit curriculum learning approach for dense prediction tasks. Our curriculum approach is defined by growing the patch size during model training, which gradually increases the task's difficulty. We integrated our curriculum into the nnU-Net framework and evaluated the methodology on all 10 tasks of the Medical Segmentation Decathlon. With our approach, we are able to substantially reduce runtime, computational costs, and CO<sub>2</sub> emissions of network training compared to classical constant patch size training. In our experiments, the curriculum approach resulted in improved convergence. We are able to outperform standard nnU-Net training, which is trained with constant patch size, in terms of Dice Score on 7 out of 10 MSD tasks while only spending roughly 50% of the original training runtime. To the best of our knowledge, our Progressive Growing of Patch Size is the first successful employment of a sample-length curriculum in the form of patch size in the field of computer vision. Our code is publicly available at <https://github.com/compai-lab/2024-miccai-fischer>.

**Keywords:** Segmentation · nnU-Net · Curriculum Learning · Resource Efficiency · Medical Segmentation Decathlon

## 1 Introduction

Automatic medical image segmentation is dominated by deep learning-based methods. Most research focuses on the development of new architectural concepts, introducing convolution-based [15,7], transformer-based [3,13] or hybrid approaches [5,4] to improve downstream segmentation performance. In contrast, few works in the medical imaging domain focus on the actual training process of networks, and models are still trained mainly by random training data sampling. Inspired by humans, Bengio et al. [2] introduced the concept of curriculum learning, teaching models to first solve simple tasks and, subsequent to that, concentrate on harder tasks. They showed that ordering samples from easy-to-hard

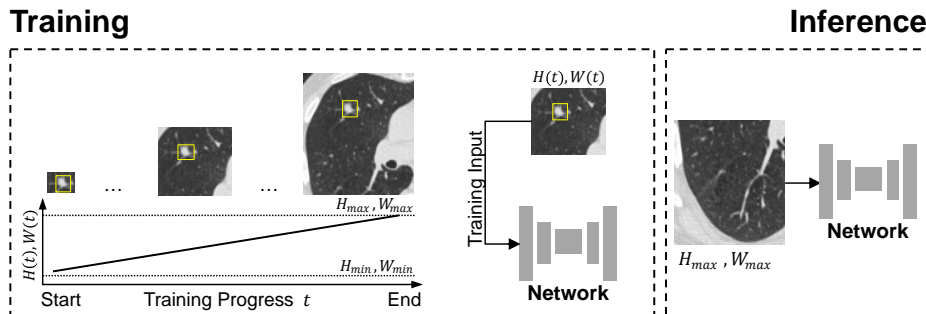
speeds up convergence and can, therefore, also improve performance compared to training with random sampling. Those training techniques can reduce the high costs of model training in terms of time, energy, and carbon footprint [17].

Different approaches have been used in the image processing field to establish such a curriculum sample ordering. Human annotations or expert knowledge can be included as problem-specific measures [2,10,19]. Task difficulty can also be directly linked to sample class membership, which was utilized for fracture classification [10]. Another frequently applied measure in the medical domain is the inter-rater expert agreement, which considerably increases annotation costs [10,19].

A more universal approach is rating the task difficulty automatically, or by incorporating processes that increase the task difficulty synthetically. Sample difficulty estimation can directly be computed from the network’s training sample loss. Jesson et al. [9] utilized hard-negative mining and an adaptive oversampling scheme for lung cancer segmentation. Missing modalities in MRI processing have been the motivation for Havei et al. [6] to train a brain tumor segmentation model that is robust against such missing data scenarios. In their curriculum learning approach, they randomly drop MRI channels with increasing probability. In the field of natural image generation, Kerras et al. [11] defined an implicit curriculum by progressively growing their image generator and discriminator layer-wise during training. With each additional layer, the output resolution is doubled, resulting in more difficult tasks. This increased convergence speed and reduced training runtime while achieving higher model performance compared to standard training. Zhao et al. [22] directly transferred this idea to semantic segmentation for their task of cervical nuclei segmentation using a U-net architecture, which grows starting from the bottleneck.

Task difficulty can also be directly defined by sample-length, which is already used in the natural language processing field in the form of sentence length [18,21,14]. In contrast, in the computer vision domain, the sample length, which would refer to patch size, has, to the best of our knowledge, not been used yet. An indirect utilization of image size is used in the Progressive Growing of GANs and its adaption to segmentation, but changes in image size are a result of the addition of network layers [11,22]. In contrast to that growing image resolution curriculum that would start segmenting a lesion from a low-resolution version of the full image, we assume that starting with segmenting a foreground object within a smaller region-of-interest patch is easier. This was our motivation to build our curriculum based on patch size instead of image resolution. In this work, we successfully apply the curriculum approach of growing patch size. Our main contributions in this work are:

1. Introduction of Progressive Growing of Patch Size curriculum for semantic segmentation and integration into the nnU-Net framework
2. Empirical verification of reduced computational costs and improved convergence compared to classical constant patch size training
3. Validation of the robustness of our approach on the 10 tasks of the Medical Segmentation Decathlon (MSD) [1]



**Fig. 1.** Proposed **Progressive Growing of Patch Size** curriculum. Our curriculum is illustrated for lung cancer segmentation with cancer bounding boxes in yellow. A fully convolutional network is able to handle inputs of different sizes. Training with our curriculum starts by training the network with minimal patch size and, with training progress, increasing the patch size dimensions linearly until the final patch size is reached. The ratio between foreground and background voxels is bigger for small patch sizes and decreases for bigger patch sizes. For inference, the maximal patch size is used.

## 2 Methods and Materials

### 2.1 Progressive Growing of Patch Size as Curriculum

We establish our curriculum by changing the patch size of the input volume. We assume training on smaller foreground patches is an easier optimization task than on larger patches. Given a fixed foreground-to-background patch ratio, the class balance is maximal for the minimal patch size and decreases with bigger patches. Larger patch sizes yield more global context, but we argue that this does not simplify the optimization, only affecting the maximal achievable performance.

While the GPU memory restricts the maximal patch size we can set, the network architecture itself restricts the minimal patch sizes that can be processed. Starting from the smallest possible patch size, we increase the patch size in the smallest possible steps to keep the task similarity between different patch sizes maximal. Each patch size is trained for the same number of iterations. During the model inference, the maximal possible patch size is used, which yields maximal global context. A sketch of our curriculum is shown in Fig. 1.

### 2.2 Implementation with nnU-Net

Fully convolutional networks are able to process inputs of varying sizes, which is the basis for our approach. For our experiments, we use the current state-of-the-art medical segmentation network nnU-Net, which follows the U-net architecture [15] and is a fully self-configuring pipeline, as detailed in [7]. The experiments were performed on the 3D patch-based segmentation version of nnU-Net (3d\_fullres), which has shown, on average, the best performances in the MSD in [8]. Models trained with the standard training of the nnU-Net are referred to

as models trained with the **Constant Patch Size Training Scheme** (CPS). For our **Progressive Growing of Patch Size Training Scheme** (PGPS), we only change the patch size compared to CPS. For the U-Net the number of downsampling operations restricts processable patch sizes, yielding a step-wise approximation of linear patch size growing. Furthermore, we also increase only one patch size dimension at a time to maximize task similarity.

### 2.3 Utilizing Smaller Patch Sizes for Bigger Batch Sizes

Training a network with a smaller patch size reduces GPU memory consumption. Thus, PGPS has lower GPU memory consumption during most of the training compared to CPS. This enables the use of larger batch sizes in the lower patch size phases of PGPS at low GPU memory costs. With a higher batch size, we expect an increase in performance and convergence speed. Besides, we force a foreground-background ratio of one, which is the default effective ratio of nnU-Net that uses a batch size of two, a result of its maximal patch size heuristic.

## 3 Experiments and Results

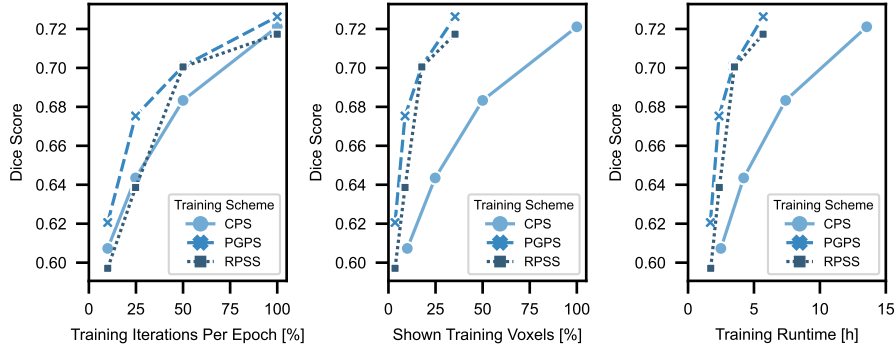
### 3.1 Lung Cancer Segmentation

**Experiments** We compare nnU-Net instances trained via CPS (default nnU-Net) and nnU-Net instances trained via our PGPS on the MSD task of lung cancer segmentation. To evaluate the effect of growing patch size, we also train models with **Random Patch Size Sampling** (RPSS). For that, at each training iteration, a random patch size of all patch sizes used for PGPS is picked, and a training iteration with the chosen patch size is performed. At testing time, all configurations use the same patch size, which is the nnU-Net default patch size.

To evaluate the convergence properties of the three different curricula, we train models for different numbers of training iterations per epoch. The standard nnU-Net 250 iterations per epoch refers to the 100% scenario. Moreover, we also train models with only 10%, 25%, 50% of the iterations per epoch. For each model training, we track the average training runtime and compute the mean Dice Score and the general number of voxels shown to the network during training.

Performance regarding training iterations, iterated voxels, and overall training runtime for trained models are plotted in Fig. 2. Exact values are given in the Supplemental Table 1. Performance of the 100% iterations per epoch training of the CPS are taken from [8].

**Results** Focusing on the number of training iterations per epoch (Fig. 2 [left]), PGPS models outperformed RPSS and CPS models in Dice Score for each trained configuration. RPSS is outperformed by CPS and PGPS for all training configurations except for 50%, where it outperforms CPS. When comparing the same models regarding training efficiency, in terms of shown voxels during training (Fig. 2 [center]) and training runtime (Fig. 2 [right]), PGPS and RPSS have



**Fig. 2.** Segmentation performance of models trained with **Progressive Growing of Patch Size (PGPS)**, **Random Patch Size Sampling (RPSS)** and **Constant Patch Size (CPS)** for different numbers of training iterations per epoch on the MSD Lung Cancer Task. Models were trained with 10%, 25%, 50%, and 100% of the training iterations per epoch, where 100% represents the default 250 iterations per epoch of nnU-Net. Dice Scores are averaged over a 5-fold cross-validation. All three plots are referring to the same models but are evaluated regarding different measures.

drastically increased efficiency. The number of shown training voxels for PGPS and RPSS is substantially reduced for each configuration to only 35.5% of the original CPS training amount. Training runtime is reflecting the trend of voxels shown. Training with 100% iterations per epoch is only running for 5.7 hours for PGPS instead of 13.55 hours for CPS.

### 3.2 Medical Segmentation Decathlon

**Experiments** To evaluate the robustness of PGPS on a variety of medical segmentation tasks, we tested our curriculum learning approach on all ten different tasks of the MSD. Detailed descriptions of the tasks are given in [1]. Various anatomies have to be segmented, including tumors, vessels, and healthy organs on single or multi-modality input, covering CT and MRI images. The MSD covers binary and multi-class segmentation tasks and dataset sample sizes ranging from 30 to 750 samples. Dice Score, runtime, and total number of voxels shown during training are computed over a 5-fold cross-validation as in [8]. Besides, we adopt *CodeCarbon* [16] to track the CO<sub>2</sub>-equivalents of our training process. Dice Scores of CPS models are taken from [8]. Because of that, we extrapolated the runtime values for CPS from the runtimes of PGPS by multiplying the averaged measured time for one epoch of the maximal patch size times the number of training epochs. Used nnU-Net model architectures and training hyperparameters can be seen for each MSD task in the Appendix of [8]. Concrete patch sizes used for PGPS are given in the Supplemental Table 2. Furthermore, we also computed paired one-sided Wilcoxon signed-rank tests to assess if the perfor-

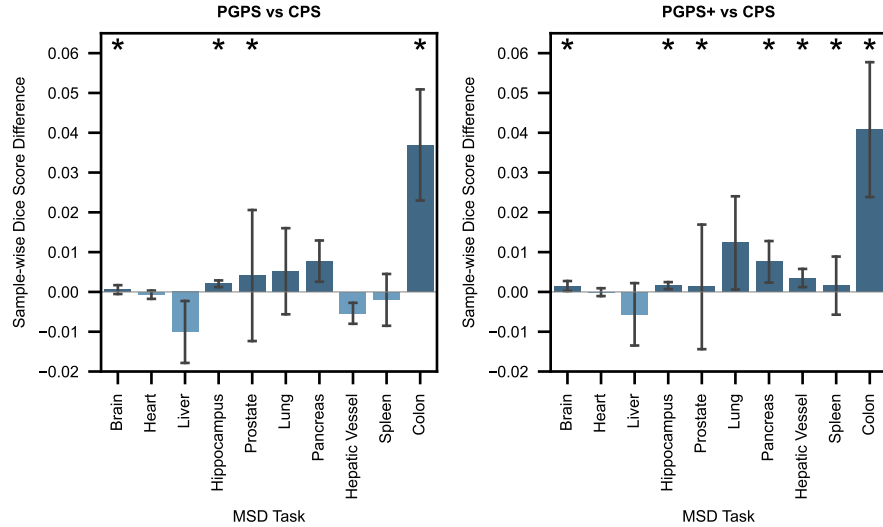
**Table 1.** Performance of **Constant Patch Size Training Scheme (CPS)**, **Progressive Growing of Patch Size Training Scheme (PGPS)**, and PGPS with increased batch size (PGPS+) on Medical Segmentation Decathlon. CPS refers to standard nnU-Net training. (**Dice Score:** Evaluated in 5-fold Cross-Validation as in [8]; **Runtime:** Training runtime per fold on one NVIDIA A100-SXM4-80GB GPU; **Voxels Shown:** total number of voxels iterated during training normalized to CPS training; **CO<sub>2</sub>-eq:** mean CO<sub>2</sub>-equivalent for training a single nnU-Net instance)

MSD Task	Scheme	Dice Score ↑	Runtime ↓	Voxels Shown ↓	CO <sub>2</sub> -eq ↓
Brain	CPS	0.7411	27.53 h	100.00 %	10.44 kg
	PGPS	0.7412	11.29 h	38.08 %	4.31 kg
	PGPS+	<b>0.7421</b>	13.83 h	51.78 %	5.21 kg
Heart	CPS	<b>0.9328</b>	15.04 h	100.00 %	6.24 kg
	PGPS	0.9321	6.44 h	35.35 %	2.64 kg
	PGPS+	0.9328	7.06 h	39.35 %	2.89 kg
Liver	CPS	<b>0.7971</b>	11.09 h	100.00 %	4.70 kg
	PGPS	0.7891	5.14 h	38.08 %	2.14 kg
	PGPS+	0.7938	6.70 h	51.78 %	2.77 kg
Hippocampus	CPS	0.8891	2.14 h	100.00 %	0.90 kg
	PGPS	<b>0.8911</b>	1.37 h	35.05 %	0.54 kg
	PGPS+	0.8907	1.42 h	35.23 %	0.56 kg
Prostate	CPS	0.7537	10.14 h	100.00 %	4.33 kg
	PGPS	<b>0.7566</b>	5.18 h	30.32 %	2.11 kg
	PGPS+	0.7531	5.75 h	36.33 %	2.31 kg
Lung	CPS	0.7211	13.55 h	100.00 %	5.59 kg
	PGPS	0.7263	5.70 h	35.35 %	2.35 kg
	PGPS+	<b>0.7333</b>	6.13 h	39.35 %	2.55 kg
Pancreas	CPS	0.6745	11.01 h	100.00 %	4.68 kg
	PGPS	<b>0.6824</b>	5.11 h	33.93 %	2.07 kg
	PGPS+	0.6822	5.19 h	35.92 %	2.30 kg
Hepatic Vessel	CPS	0.6837	15.12 h	100.00 %	6.06 kg
	PGPS	0.6782	6.18 h	36.19 %	2.51 kg
	PGPS+	<b>0.6871</b>	6.65 h	40.36 %	2.72 kg
Spleen	CPS	0.9638	11.24 h	100.00 %	4.76 kg
	PGPS	0.9621	5.22 h	36.30 %	2.15 kg
	PGPS+	<b>0.9654</b>	5.37 h	45.81 %	2.31 kg
Colon	CPS	0.4553	10.94 h	100.00 %	3.77 kg
	PGPS	0.4925	4.36 h	32.21 %	1.69 kg
	PGPS+	<b>0.4967</b>	4.50 h	34.55 %	1.77 kg

mance differs significantly (P-value < 0.05). For that, we pair validation Dice Scores of each sample over all cross-validation splits between CPS and PGPS.

We repeat the experiments with increased batch size for the PGPS curriculum. In our approach, we increase the batch size of the current patch size stage to the extent that the total number of voxels in an input tensor is equal to or lower than that of the following patch size stage. In this way, we avoid drops in the number of input voxels per iteration. We refer to this adaption of PGPS as PGPS+ in the later text. The concrete patch and batch sizes used for PGPS and PGPS+ are given in the Supplemental Table 2.

**Results** Results on the MSD tasks are given in Table 1. In Fig. 3, we plotted the sample-wise Dice Score differences between our new training curricula and standard CPS. On averaged Dice Scores over the 5-fold cross-validation, the PGPS models outperform CPS models in 6 out of 10 tasks. When PGPS outperformed CPS, it achieved an average fold-wise increase of  $1.8\% \pm 2.9\%$  in Dice



**Fig. 3.** Performance of **Progressive Growing of Patch Size** (PGPS), PGPS with increased batch size (PGPS+) and **Constant Patch Size** (CPS) on the MSD tasks. Here, we show the average sample-wise difference in Dice Score and its standard error over all samples (combination over all five validation folds). A positive difference represents wins of the PGPS curricula over the CPS training (here marked in dark blue). On this measure, PGPS outperforms CPS in 6 of 10 tasks, and PGPS+ outperforms CPS in 8 of 10 tasks. Significant difference is tested via a one-sided Wilcoxon signed-rank tests ( $P$ -value  $< 0.05$ ). We observed only significant improvements, while decreases were not significant.

Score compared to CPS, while for tasks CPS outperformed PGPS, there was only an average fold-wise difference of  $0.5\% \pm 0.4\%$  in performance. Comparing the performance differences regarding each sample, significance testing resulted in 4 significant improvements and no significant performance degradation. Full training of the nnU-Net with our PGPS curriculum reduces the runtime on average to  $46.09\% \pm 6.8\%$  compared to CPS (default nnU-Net training) by only iterating  $35.09\% \pm 2.30\%$  of the CPS' seen voxels and reducing CO<sub>2</sub> equivalents to  $48.22\% \pm 9.34\%$ .

With the additional adjustment of the batch size for PGPS+, we are able to outperform CPS in 7 out of 10 MSD tasks on the mean of fold-wise Dice Scores. When PGPS+ outperformed CPS, it achieved an average increase of  $1.7\% \pm 3.1\%$  in fold-wise average Dice Score compared to CPS, while when it lost, there was only an average difference of  $0.2\% \pm 0.2\%$ . Comparing the performance differences between each sample, our PGPS+ even outperforms CPS in 8 out of 10 tasks. Significance testing resulted in 7 significant improvements and no significant performance degradation. The model training only takes up

an average of  $50.59\% \pm 7.6\%$  of the original time of CPS, while iterated voxels are reduced to  $41.05\% \pm 6.20\%$  and CO<sub>2</sub> equivalents to  $50.59\% \pm 5.55\%$ .

## 4 Discussion and Conclusion

We introduced a novel curriculum learning approach based on increasing the patch size during training and compared it with constant patch size training.

Our curriculum approach showed improved convergence in the experiments, only significant improvements and no significant decreases in performance occurred. Especially, our curriculum outperforms classical CPS on the MSD most difficult task of colon cancer by a large margin. Better convergence is a proven property of optimal curricula [20]. In the lung cancer experiment, we showed that the performance gain per training voxel is drastically increased compared to CPS. This was observed for all different training configurations, differing in the length of network training. The beneficial impact of our curriculum is based on the ordering of patches by growing size, which we assume results in an efficient, easy-to-hard order. Furthermore, PGPS also outperforms random ordering (RPSS), validating the impact of the patch size ordering. We hypothesize that with growing patch size, we ask the network subsequently more difficult questions. In contrast, Li et al. [12] used only two different patch sizes during training as a comparison method to their proposed curriculum. Their results showed a performance drop of their *patch-to-whole* training over CPS training. We assume that our approach benefits from maximal patch size stages.

Moreover, our curriculum can utilize larger batch sizes within the same GPU hardware restrictions. In our experiments, this resulted in a performance improvement of PGPS+ over PGPS, as higher batch sizes generally can lead to better convergence and performance.

Besides the improved convergence, our proposed curriculum also drastically reduced the training runtime and CO<sub>2</sub> emissions to roughly half of the original values for PGPS and PGPS+. This is a result of the cheaper operations due to smaller input sizes. Minimal training runtime is defined by the total number of seen voxels during training and depends highly on the technical implementation.

We acknowledge that this is a proof of concept work, relying on the hyperparameters of nnU-Net. More experiments are needed to explore the different effects of using another minimal patch size, number of patch size stages, batch size, normalization, and training scheduler. Furthermore, our approach is not limited to fully convolutional networks but could also be applied to transformers [3,13] and hybrid model architectures [4,5] as well as object detection tasks.

Overall, the proposed Progressive Growing of Patch Size is a resource-efficient curriculum strategy that drastically reduces the training runtime and furthermore can also lead to improved convergence. The performance gain per voxel is substantially higher than for classical constant patch size training, supporting our hypothesis of better convergence. Improving the efficiency of network training is particularly important considering the growing carbon footprint of training deep learning models [17].



**Acknowledgments.** Stefan Fischer has received funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 515279324 / SPP 2177. Johannes Kiechle was supported by the DAAD programme Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research. We thank the team of DKFZ Medical Image Computing research group for providing nnU-Net results for in-depth performance comparison.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., et al.: The medical segmentation decathlon. *Nature Communications* **13**(1), 4128 (2022). <https://doi.org/10.1038/s41467-022-30695-9>
2. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: *Proceedings of the International Conference on Machine Learning*. pp. 41–48. PMLR (2009)
3. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
4. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images. In: *Brainlesion Workshop at International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 272–284. Springer (2021)
5. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: UNETR: Transformers for 3D medical image segmentation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 574–584 (2022)
6. Havaei, M., Guizard, N., Chapados, N., Bengio, Y.: HeMIS: Hetero-modal image segmentation. In: *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 469–477. Springer (2016)
7. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
8. Isensee, F., Jäger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: Automated design of deep learning methods for biomedical image segmentation. *arXiv preprint arXiv:1904.08128* (2019)
9. Jesson, A., Guizard, N., Ghahlehjeh, S.H., Goblot, D., Soudan, F., Chapados, N.: Cased: Curriculum adaptive sampling for extreme data imbalance. In: *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 639–646. Springer (2017)
10. Jiménez-Sánchez, A., Mateus, D., Kirchoff, S., Kirchoff, C., Biberthaler, P., Navab, N., González Ballester, M.A., Piella, G.: Medical-based deep curriculum learning for improved fracture classification. In: *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 694–702. Springer (2019)

11. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017)
12. Li, H., Liu, X., Boumaraf, S., Liu, W., Gong, X., Ma, X.: A new three-stage curriculum learning approach for deep network based liver tumor segmentation. In: Proceedings of International Joint Conference on Neural Networks. pp. 1–6. IEEE (2020)
13. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
14. Platanios, E.A., Stretcu, O., Neubig, G., Poczos, B., Mitchell, T.M.: Competence-based curriculum learning for neural machine translation. arXiv preprint arXiv:1903.09848 (2019)
15. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 234–241. Springer (2015)
16. Schmidt, V., Goyal, K., Joshi, A., Feld, B., Conell, L., Laskaris, N., Blank, D., Wilson, J., Friedler, S., Luccioni, S.: CodeCarbon: Estimate and track carbon emissions from machine learning computing (2021). <https://doi.org/10.5281/zenodo.4658424>, v2.3.4
17. Selvan, R., Bhagwat, N., Wolff Anthony, L.F., Kanding, B., Dam, E.B.: Carbon footprint of selecting and training deep learning models for medical image analysis. In: Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 506–516. Springer (2022)
18. Spitkovsky, V.I., Alshawi, H., Jurafsky, D.: From baby steps to leapfrog: How “less is more” in unsupervised dependency parsing. In: Proceedings of Annual Conference of the North American Chapter of the Association for Computational Linguistics. pp. 751–759. Association for Computational Linguistics (2010)
19. Wei, J., Suriawinata, A., Ren, B., Liu, X., Lisovsky, M., Vaickus, L., Brown, C., Baker, M., Nasir-Moin, M., Tomita, N., Torresani, L., Wei, J., Hassanpour, S.: Learn like a pathologist: Curriculum learning by annotator agreement for histopathology image classification. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2473–2483 (2021)
20. Weinshall, D., Cohen, G., Amir, D.: Curriculum learning by transfer learning: Theory and experiments with deep networks. In: Proceedings of International Conference on Machine Learning. pp. 5238–5246. PMLR (2018)
21. Zaremba, W., Sutskever, I.: Learning to execute. arXiv preprint arXiv:1410.4615 (2014)
22. Zhao, J., Dai, L., Zhang, M., Yu, F., Li, M., Li, H., Wang, W., Zhang, L.: PGU-net+: Progressive growing of U-net+ for automated cervical nuclei segmentation. In: Multiscale Multimodal Medical Imaging Workshop at International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 51–58. Springer (2020)