

Mixed-Block Neural Architecture Search for Medical Image Segmentation

Bosma, Martijn M.A.; Dushatskiy, Arkadiy; Grewal, Monika; Alderliesten, Tanja; Bosman, Peter A.N.

DOI

[10.1117/12.2611428](https://doi.org/10.1117/12.2611428)

Publication date

2022

Document Version

Final published version

Published in

Medical Imaging 2022

Citation (APA)

Bosma, M. M. A., Dushatskiy, A., Grewal, M., Alderliesten, T., & Bosman, P. A. N. (2022). Mixed-Block Neural Architecture Search for Medical Image Segmentation. In O. Colliot, I. Isgum, B. A. Landman, & M. H. Loew (Eds.), *Medical Imaging 2022: Image Processing* Article 120320S (Progress in Biomedical Optics and Imaging - Proceedings of SPIE; Vol. 12032). SPIE. <https://doi.org/10.1117/12.2611428>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://spiedigitallibrary.org/conference-proceedings-of-spie)

Mixed-block neural architecture search for medical image segmentation

Martijn M. A. Bosma, Arkadiy Dushatskiy, Monika Grewal, Tanja Alderliesten, Peter Bosman

Martijn M. A. Bosma, Arkadiy Dushatskiy, Monika Grewal, Tanja Alderliesten, Peter Bosman, "Mixed-block neural architecture search for medical image segmentation," Proc. SPIE 12032, Medical Imaging 2022: Image Processing, 120320S (4 April 2022); doi: 10.1117/12.2611428

SPIE.

Event: SPIE Medical Imaging, 2022, San Diego, California, United States

Mixed-Block Neural Architecture Search for Medical Image Segmentation

Martijn M.A. Bosma^{a, b}, Arkadiy Dushatskiy^{*a}, Monika Grewal^{*a}, Tanja Alderliesten^c, and Peter A. N. Bosman^{a, b}

^aCentrum Wiskunde & Informatica, Science Park 123, 1098XG, Amsterdam, the Netherlands

^bDelft University of Technology, Mekelweg 5, 2628CD, Delft, the Netherlands

^cLeiden University Medical Center, Albinusdreef 2, 2333ZA, Leiden, the Netherlands

ABSTRACT

Deep Neural Networks (DNNs) have the potential for making various clinical procedures more time-efficient by automating medical image segmentation. Due to their strong, in some cases human-level, performance, they have become the standard approach in this field. The design of the best possible medical image segmentation DNNs, however, is task-specific. Neural Architecture Search (NAS), i.e., the automation of neural network design, has been shown to have the capability to outperform manually designed networks for various tasks. However, the existing NAS methods for medical image segmentation have explored a quite limited range of types of DNN architectures that can be discovered. In this work, we propose a novel NAS search space for medical image segmentation networks. This search space combines the strength of a generalised encoder-decoder structure, well known from U-Net, with network blocks that have proven to have a strong performance in image classification tasks. The search is performed by looking for the best topology of multiple cells simultaneously with the configuration of each cell within, allowing for interactions between topology and cell-level attributes. From experiments on two publicly available datasets, we find that the networks discovered by our proposed NAS method have better performance than well-known handcrafted segmentation networks, and outperform networks found with other NAS approaches that perform only topology search, and topology-level search followed by cell-level search.

1. INTRODUCTION

A growing amount of clinical applications, such as computer-aided diagnostic systems, are benefiting from recent advances in automated medical image segmentation, most notably from Deep Neural Networks (DNNs).^{1,2} Given a medical scan, a DNN can provide contours of organs or regions of interest (e.g., tumors), with clinically acceptable segmentation quality within a matter of seconds.³ Designing a State-of-the-Art (SotA) DNN, however, is often task-specific. In order to design a DNN, choices have to be made for the topology of the network such as depth, and connections between cells (a cell is a group of operations that transform the feature maps in a DNN), as well as the configuration of each cell, e.g., convolutional kernel size, or activation function. This gives rise to an inconceivably large amount of network architecture design possibilities, which is impossible to manually navigate through in an exhaustive fashion, or even by means of intelligent design, while ensuring the best choices are made.

Neural Architecture Search (NAS), i.e., the automated design of neural network architectures, can effectively and efficiently search through this space of possible network architecture designs and find a network that is highly tailored to the task at hand.⁴ While research on NAS for medical image segmentation has not been as elaborate as for natural image classification, it has already shown promising results by outperforming the SotA architectures.⁵⁻⁷ In our opinion, further research on NAS for medical image segmentation can make its contributions even more significant.

NAS involves three key components: (1) The search space (the set of all possible networks given the specified architectural constraints); (2) the search algorithm (the algorithm to navigate the search space); (3) performance

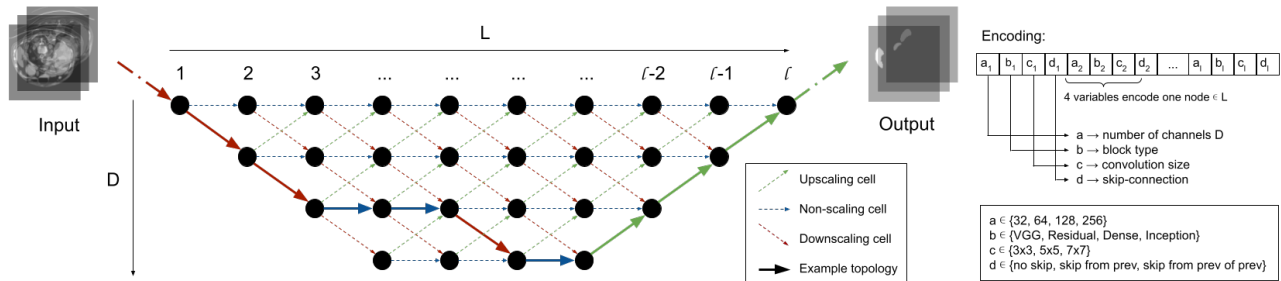
*equal contribution

estimation (the choices made to score a network’s performance, such that these networks can be ranked by the search algorithm). So far, NAS research has been more elaborate for image classification tasks.⁴ Potentially, strong performing search algorithms, and fast and accurate performance estimation methods can be translated to the medical image segmentation domain. The search space, however, is specific to segmentation tasks. Several papers have proposed search spaces for medical image segmentation. These include searching for a U-Net like encoder-decoder topology i.e., by adapting only the cells within,⁷ by searching for the best topology followed by the best convolution size within cells,⁵ or a combination of topology and cell search using continuous relaxation and a so-called Super-network.⁶ With each of these NAS search spaces it has been possible to find networks that perform slightly better than a standard U-Net. A semi-automatic network architecture design and training method known as nnU-Net,⁸ however, is still considered SotA for many image segmentation tasks.

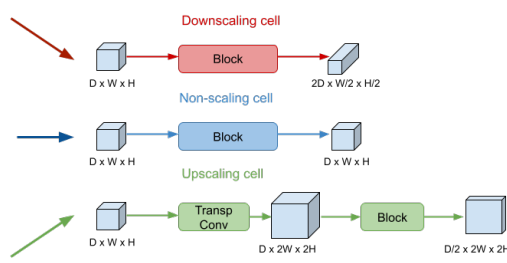
We believe that within the scope of a U-Net-like topology search, there is a room for making the search space more flexible – instead of repeating the same cell structure, we propose to allow various cell structures in a network, potentially resulting in a better performance. Further, we believe that the configuration of cells can benefit from existing knowledge about networks with a strong performance on classification tasks. Therefore, we propose to search through a pre-selected pool of configurations (which are taken from advanced well-known classification networks^{9–11}) instead of searching for a cell configuration from scratch. Apart from benefiting from advanced performance of these well known classification networks, this proposed improvement will also help avoiding the explosive growth of the search space caused by searching the configuration of each cell from scratch. Finally, in contrast to recent research,⁵ where topology level search is followed by cell level search, we search for both topology as well as the configuration of each cell simultaneously. This allows to take into consideration the possible interaction between the topology-level search and cell-level, potentially yielding better performing networks. The combination of these improvements results in the proposed approach, which we will further refer to as Mixed-Block NAS (MB-NAS).

2. METHOD

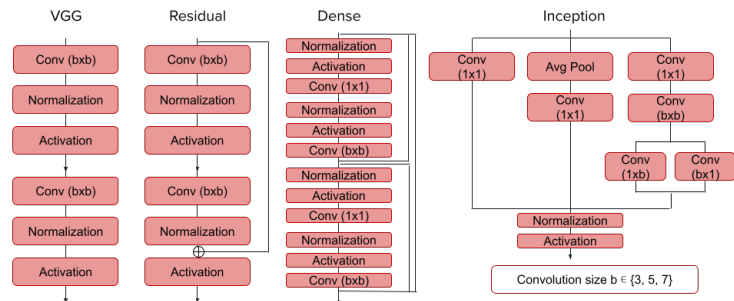
2.1 Search space



(a) Topology-level search and encoding.



(b) Cell types.



(c) Block types.

Figure 1: Description of MB-NAS search space.

At **topology-level** (Figure 1a), the search space contains all possible network architectures resulting from varying connections between cells of different types. We consider three possibilities for a cell (Figure 1b): a downscaling, an upscaling, or a non-scaling cell. A downscaling cell means that the input image resolution is halved, while the number of channels is doubled. Upscaling is the opposite operation, i.e., doubling the resolution and halving the number of channels, while a non-scaling cell changes neither the resolution nor the number of channels in the feature maps. The input to a cell can be either just the feature maps from the preceding cell (no skip-connection), or the feature maps of the preceding cells concatenated with the feature maps from any previous cell with the same feature maps spatial dimensions (a single skip-connection).

At **cell-level**, different configurations of each cell are searched. The configuration of a cell is encoded by two variables: the type of block (a block is an organised structure consisting of multiple convolution and normalisation layers, as well as activation functions) within the cell, and the convolutional kernel size within the block. Instead of searching for the topology of the blocks within a cell, which would make the search space incredibly large, we used predefined blocks derived from previously SotA architectures for classification (Figure 1c). In this work we consider VGG blocks¹² which are standard in U-Net, as well as Residual blocks,⁹ Dense blocks¹⁰ and Inception blocks.¹¹ In this way, we allow the search space to choose a different cell configuration at every edge (instead of repeating one throughout the network) while preventing the further growth of the search space.

The network architecture is represented by connections between a fixed number of nodes (fixed to 10 in our experiments). Each node l (Figure 1a) is represented by 4 categorical variables: a_l = number of channels, b_l = the block type, c_l = convolution size, and d_l = skip-connection source. The cell type is derived by the difference in number of channels between two nodes. The topology of the neural network is encoded by variables a_l and d_l at each node. Variable a_l is restrained to only be allowed to take on double, equal or half the number of channels contained in the previous cell. The cell-level search is represented by variables b_l and c_l at each node. Note that the standard U-Net shape is included in the topology-level search space. The resulting search space contains $1.14 * 10^{18}$ possible networks.

2.2 Datasets

In our experiments, we used two datasets that can be found on the Medical Decathlon¹³ challenge website[†]. The first is a collection of 3D MRI-scans of the prostate, alongside their multi-class (Central and Peripheral prostate zones) segmentations. It contains 32 multi-modal scans, represented as a collection of 602 2D slices total. The scans were normalised, using Z-normalisation. The second contains 3D-CT scans and single class segmentations of the spleen. It contains 41 single modality scans represented as 3650 2D slices. The scans were clipped and normalised.

2.3 Implementation details

We use Local Search (LS) as the search algorithm as it has been shown to be a strong baseline for NAS.¹⁴ Specifically, we use a first-improvement approach with a variable neighbourhood of 1. This means that for each variable (iterated in a random order), all possible options are probed, and the best performing is chosen. Once all variables are considered, we start over, until no improvements can be found anywhere. Due to limited computational resources, the NAS was run for only 150 network evaluations. Each network was evaluated using the validation Dice score averaged over the last 20% of training epochs. The number of epochs (100 epochs for Prostate, 50 for Spleen) was decided such that saturation was ensured based on preliminary runs. We used the average score of 5-fold patient-level cross-validation, repeated for 3 different network initialisations, as performance estimator. Multiple initialisations reduce the noise in scoring, providing more reliable information to the search algorithm, making it easier to navigate the search space. The use of multiple folds decreases noise caused by data splits.

For network training, we used the ADAM optimiser with a learning rate 10^{-3} and polynomial decay with an exponent of 0.9. The loss function was foreground soft Dice. The batch size was 32 and input image size was 128x128. The data was augmented using scaling, shifting, rotating, flipping, and brightness adjustment to avoid overfitting.

[†]<http://medicaldecathlon.com/>

2.4 Experimental setup

We performed experiments to compare the proposed search space design (Mixed-Block, or **MB-NAS**) against the following alternative approaches. In the first considered alternative search space (**Macro-NAS**) only topology search is performed. All block types are fixed to be standard U-Net blocks (VGG). The second considered search space (**Micro-NAS**) has the U-Net topology, where only the block type and convolution sizes are subject to the search. In the third alternative search space (**Bilevel-NAS**), a bi-level approach was used which means that first the topology-level is searched, and then the convolution size, similar to the Coarse-to-Fine approach.⁵

The best networks found by the proposed search space were evaluated against two hand-crafted neural network architectures: standard **U-Net**, and U-Net with a ResNet-50 encoder (**ResU-Net**). Implementations are taken from the Pytorch Segmentation Models library.¹⁵

3. RESULTS AND DISCUSSION

The performance of the different search spaces (**MB-NAS**, **Macro-NAS**, **Micro-NAS**, **Bilevel-NAS**) can be seen in Figure 2a and 2b. The the architectures of the best networks from MB-NAS are visualised in Figure 2c. The performance of the best networks from different NAS approaches, **U-Net**, and **ResU-Net** is summarised in Table 1.

Table 1: Performance values of U-Net, ResU-Net, and the best networks from different NAS approaches. The values are averaged over 5 seeds on a 5-fold cross-validation. Standard deviations are calculated based on single seed scores on all 5-folds. Best values in each column are highlighted in bold. DSC: Dice-Sorensen similarity coefficient, HD: Hausdorff distance (95% cutoff), SD: Surface Dice (2mm threshold), MMAC: Mega Multiply-ACcumulate operations, #Params: number of parameters, $\times 10^6$.

| Model | Prostate dataset | | | | |
|--------------------|--------------------------------------|-------------------------------------|--------------------------------------|------------|-------------|
| | DSC | HD | SD | MMAC | #Params |
| U-Net | 0.6702 \pm 0.004 | 8.833 \pm 0.348 | 0.6046 \pm 0.006 | 302 | 18.4 |
| ResU-Net | 0.6580 \pm 0.004 | 9.441 \pm 0.325 | 0.5705 \pm 0.006 | 166 | 32.5 |
| Macro-NAS | 0.6593 \pm 0.004 | 8.606 \pm 0.470 | 0.5977 \pm 0.008 | 256 | 3.39 |
| Micro-NAS | 0.6796 \pm 0.007 | 8.394 \pm 0.251 | 0.6203 \pm 0.008 | 1,295 | 22.7 |
| Bilevel-NAS | 0.6702 \pm 0.005 | 8.492 \pm 0.541 | 0.6134 \pm 0.008 | 414 | 6.77 |
| MB-NAS | 0.6760 \pm 0.010 | 8.419 \pm 0.324 | 0.6192 \pm 0.004 | 644 | 3.04 |

| Model | Spleen dataset | | | | |
|--------------------|--------------------------------------|-------------------------------------|--------------------------------------|------------|-------------|
| | DSC | HD | SD | MMAC | #Params |
| U-Net | 0.9578 \pm 0.002 | 1.412 \pm 0.030 | 0.917 \pm 0.002 | 302 | 18.4 |
| ResU-Net | 0.9464 \pm 0.004 | 1.625 \pm 0.042 | 0.905 \pm 0.003 | 166 | 32.5 |
| Macro-NAS | 0.9566 \pm 0.001 | 1.467 \pm 0.063 | 0.9167 \pm 0.001 | 255 | 3.39 |
| Micro-NAS | 0.9567 \pm 0.002 | 1.388 \pm 0.026 | 0.9177 \pm 0.004 | 795 | 2.83 |
| Bilevel-NAS | 0.9553 \pm 0.001 | 1.449 \pm 0.069 | 0.9145 \pm 0.005 | 415 | 6.78 |
| MB-NAS | 0.9592 \pm 0.002 | 1.385 \pm 0.036 | 0.9189 \pm 0.002 | 1,294 | 22.7 |

3.1 Spleen dataset

The best performing network on the Spleen dataset was found using the proposed search space, MB-NAS. In Table 1, one can see that this network shows the best performance by all three considered performance metrics. Note that the topology (Figure 2c, top row) is quite different from the standard U-Net. The network is shallower, potentially indicating that smaller, more intricate features, are more important in the segmentation of the spleen, than larger, more elaborate ones. This could be due to the small size of the organ in relation to the image. Furthermore, blocks of all four types, are included in the architecture, as well as different convolution sizes. This means that both the topology-level and cell-level search space parts were utilised to find this network; an indication towards the performance increase from the designed search space.

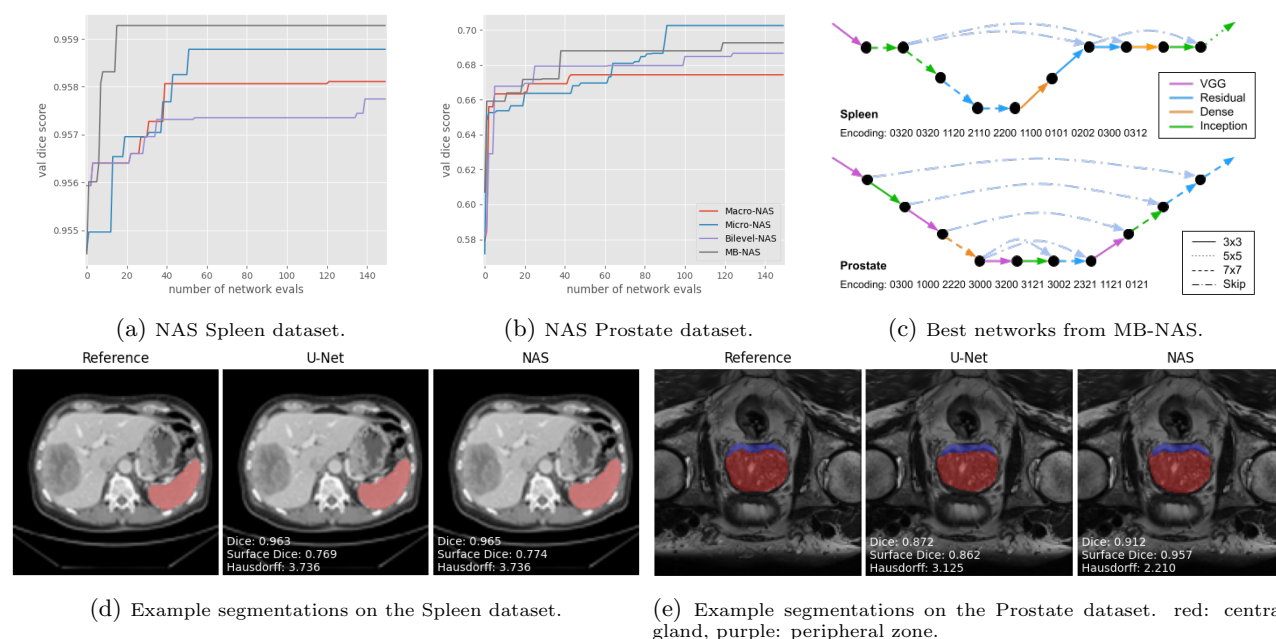


Figure 2: The progress of the LS algorithm for the (a) Spleen and (b) Prostate datasets, respectively. The best networks found by LS with MB-NAS are shown in (c). (d) and (e) show example segmentations.

3.2 Prostate dataset

Different than for the Spleen dataset, the best network (Figure 2c, bottom row) is very similar to the U-Net architecture. This provides additional evidence to the argument that the architecture of the best DNN is task-specific. This also indicates that the U-Net topology is best suited for the underlying task, giving an added advantage to Micro-NAS, wherein different blocks and convolution sizes are searched within the U-Net topology. Consequently, the use of Micro-NAS yields the best performance. However, it should be noted that using MB-NAS, we could find a network with comparable performance despite the fact that the search was performed for both topology as well as cell configuration. It is worth noting that no block seems to be ultimately preferable to other blocks, i.e. outperforms other blocks at every location of the network. This indicates the advantage of searching for a mixed-block configuration such that different blocks can be used throughout the network. Additionally, the differences in the validation DSC of poorly performing networks compared to the best networks is much larger for the Prostate dataset, indicating a more difficult task. Further, Table 1 shows that the best networks from NAS outperform the manually hand-crafted U-Net and ResU-Net networks.

3.3 Comparison with SotA

It can be argued that the results for the Prostate dataset in Table 1 are not at the same level as the SotA results given by e.g., nnU-Net.⁸ This is due to higher image resolution used by nnU-Net, additional preprocessing, carefully chosen data augmentations, post-processing, and an advanced inference method, that are all different from the setup used in NAS in this paper. Therefore, the network performance when using the nnU-Net training and evaluation setup, was also evaluated with the found networks from NAS. For Prostate, the 5-fold cross-validation Dice-Sorensen coefficient is **0.7325** for MB-NAS vs. **0.7315** for nnU-Net. It should be noted that the comparison is done between the network found by NAS, trained in the nnU-Net environment, and a U-Net architecture that is tailored to the data according to the nnU-Net heuristics. The architectures found by NAS were found using entirely different settings. This gives an unfair advantage to nnU-Net. Nevertheless, the architecture found by NAS still performs slightly better to nnU-Net in these settings, which is remarkable. For the Spleen dataset, the performance level is quite similar to the results reported by nnU-Net. Nevertheless,

the network was validated in this environment as well. The 5-fold cross-validation Dice-Sorensen coefficient is **0.9467** for MB-NAS vs. **0.9466** for nnU-Net, indicating equal performance.

3.4 Statistical significance

A Friedman test was conducted on the different network performances to compare the effect of the architecture on the segmentation accuracy. The test was performed on individual scores for 5 differently seeded network initialisations on all 5 folds of a datasplit. The Friedman test reveals a statistically significant difference between network performances from Table 1 for the Dice-Sorensen similarity coefficient (Spleen dataset : ($\chi^2(5) = 49.14, p < 0.0001$); Prostate dataset : ($\chi^2(5) = 16.98, p < 0.005$)). A post-hoc analysis, using the Wilcoxon signed-rank test, found that the difference in performance of the Dice-Sorensen similarity coefficient was statistically significant between the best network found by NAS and U-Net for the Spleen dataset ($Z = 48.0, p < 0.05$). For the Prostate dataset the difference was not significant.

3.5 Computational limits and possible extensions

We note that due to limited computational resources, only one run of NAS was performed per search space, and for a limited number of network evaluations. Longer experiments with multiple runs would have helped draw more definitive conclusions. Also, running NAS experiments for more datasets may provide further insights about the specific characteristics of network architecture design required for good performance across datasets. Additionally, the search space could be further extended by adding other block types, like Image Transformer¹⁶ blocks.

4. CONCLUSIONS

We have proposed an approach for medical image segmentation Neural Architecture Search, which involves a novel search space and simultaneous topology- and cell-level search strategy. In the cell-level search, we used existing knowledge from networks with high performance in image classification tasks, i.e. ResNet, DenseNet and InceptionNet, to create a pool of possible block configurations. The experiments in this paper show the added value of this approach. Overall, the results indicate that further research into search space refinement, allowing to exploit key features of what accounts for good deep learning performance, may yet push the boundaries of what can be achieved with deep neural networks for medical image processing.

REFERENCES

- [1] Bernard, O. et al., “Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis,” *IEEE Transactions on Medical Imaging* **37**(11), 2514–2525 (2018).
- [2] Fauw, D. et al., “Clinically applicable deep learning for diagnosis and referral in retinal disease,” *Nature Medicine* **24**(9), 1342–1350 (2018).
- [3] Nikolov, S., Blackwell, S., Zverovitch, A., Mendes, R., Livne, M., Fauw, J. D., Patel, Y., Meyer, C., Askham, H., Romera-Paredes, B., Kelly, C., Karthikesalingam, A., Chu, C., Carnell, D., Boon, C., D’Souza, D., Moineddin, S. A., Garie, B., McQuinlan, Y., Ireland, S., Hampton, K., Fuller, K., Montgomery, H., Rees, G., Suleyman, M., Back, T., Hughes, C., Ledsam, J. R., and Ronneberger, O., “Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy,” (2021).
- [4] Elsken, T., Metzen, J. H., and Hutter, F., “Neural architecture search: A survey,” *The Journal of Machine Learning Research* **20**(1), 1997–2017 (2019).
- [5] Yu, Q., Yang, D., Roth, H., Bai, Y., Zhang, Y., Yuille, A. L., and Xu, D., “C2fnas: Coarse-to-fine neural architecture search for 3d medical image segmentation,” (2020).
- [6] Yan, X. et al., “Ms-nas: Multi-scale neural architecture search for medical image segmentation,” in [*International Conference on Medical Image Computing and Computer-Assisted Intervention*], 388–397, Springer (2020).
- [7] Weng, Y., Zhou, T., Li, Y., and Qiu, X., “Nas-unet: Neural architecture search for medical image segmentation,” *IEEE Access* **7**, 44247–44257 (2019).

- [8] Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J., and Maier-Hein, K. H., “nnu-net: A self-configuring method for deep learning-based biomedical image segmentation,” *Nature methods* **18**(2), 203–211 (2021).
- [9] He, K. et al., “Deep residual learning for image recognition,” in [*Proceedings of the IEEE conference on computer vision and pattern recognition*], 770–778 (2016).
- [10] Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K. Q., “Densely connected convolutional networks,” (2018).
- [11] Szegedy, C. et al., “Inception-v4, inception-resnet and the impact of residual connections on learning,” in [*Thirty-first AAAI Conference on Artificial Intelligence*], (2017).
- [12] Simonyan, K. and Zisserman, A., “Very deep convolutional networks for large-scale image recognition,” (2015).
- [13] Antonelli, M. et al., “The medical segmentation decathlon,” (2021).
- [14] Ottelander, T. D., Dushatskiy, A., Virgolin, M., and Bosman, P. A. N., “Local search is a remarkably strong baseline for neural architecture search,” (2020).
- [15] Yakubovskiy, P., “Segmentation models.” https://github.com/qubvel/segmentation_models (2019).
- [16] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N., “An image is worth 16x16 words: Transformers for image recognition at scale,” (2021).