DOI 10.31891/2307-5732-2024-339-4-56
УДК 004.8, 004.67, 004.42

ЛУКАЩУК БОГДАН
Національний лісотехнічний університет України
https://orcid.org/0009-0002-2365-4923
e-mail: bohdan.lukashchuk@gmail.com

СЕГМЕНТАЦІЯ ЗОБРАЖЕНЬ РАН НА ШКІРІ ЛЮДИНИ ЗА ДОПОМОГОЮ ЗГОРТКОВИХ НЕЙРОННИХ МЕРЕЖ

Це дослідження спрямоване на підвищення точності та ефективності сегментації зображень ран на шкірі людини за допомогою згорткових нейронних мереж (ЗНМ). Точна сегментація ран є важливою для ефективного лікування та моніторингу в медицині. Ми оцінили та порівняли сучасні архітектури ЗНМ, зосередившись на U-Net та моделях, оптимізованих для використання на мобільних пристроях, таких як DeeplabV3 та LRASPP.

Ми використовували набір даних Foot Ulcer Segmentation Challenge у різних умовах, включаючи попередньо навчені моделі та моделі з випадковими вагами, із різною стандартизацією вхідних даних та без неї. Бібліотека PyTorch та хмарне середовище Google Colaboratory PRO були використані для розробки, навчання, валідацію та тестування моделей.

Результати показали, що U-Net стабільно досягає високих значень індексу Дайса, особливо при використанні попередньо навчених ваг зі стандартизацією даних із найвищим значенням 0,875. LRASPP також показала хороші результати, особливо при використанні попередньо-навчених ваг, але з дещо нижчими оцінками метрики Дайса, ніж U-Net. DeeplabV3 значно відстала від інших моделей за точністю сегментації. Попередньо навчені моделі, стандартизовані на основі середнього значення та стандартного відхилення нового набору даних, показали кращі результати, ніж моделі, стандартизовані на основі статистики попереднього набору даних.

Аналіз кривих функцій втрат при навчанні та валідації виявив потенційне перенавчання моделей U-Net та DeeplabV3, про що свідчать флуктуації функцій втрат на валідаційному наборі даних. Ми пропонуємо інтегрувати методи регуляризації та стратегії доповнення даних для вирішення цієї проблеми та покращення узагальнення. Крім того, налаштування гіперпараметрів моделей та використання шарів BatchNormalization для U-Net може допомогти зменшити перенавчання.

Результати показують, що U-Net є хорошим вибором для задач сегментації ран, особливо за умови попереднього навчання і належної стандартизації вхідних даних. Для застосування на мобільних пристроях LRASPP демонструє значний потенціал. Подальша робота повинна вдосконалити ці моделі для мобільного розгортання, забезпечуючи надійні та ефективні результати.

Ключові слова: аналіз зображень, глибоке навчання, згорткові нейронні мережі, семантична сегментація, сегментація ран, аналіз медичних зображень, U-Net, DeeplabV3, LRASPP, PyTorch.

LUKASHCHUK BOHDAN
Ukrainian National Forestry University

HUMAN SKIN WOUND IMAGE SEGMENTATION USING CONVOLUTIONAL NEURAL NETWORKS

This study aims to enhance human skin wound segmentation accuracy and efficiency using convolutional neural networks (CNNs). Accurate wound segmentation is crucial for effective treatment and monitoring in medical applications. We evaluated and compared modern CNN architectures, focusing on U-Net, DeeplabV3, and LRASPP across different setups, to identify models suitable for mobile deployment with efficient computation and minimal resource use.

Our experiments utilized the Foot Ulcer Segmentation Challenge dataset across various setups, including pre-trained and non-pre-trained models, with and without data standardization. We used PyTorch library for structured experiment management and conducted training and validation on Google Colaboratory PRO.

Results showed that U-Net consistently achieved high Dice scores, particularly in pre-trained models with dataset standardization, with the highest score of 0.875. LRASPP also performed well, especially in pre-trained setups, with slightly lower Dice scores than U-Net. DeeplabV3 lagged behind the other models in segmentation accuracy. Notably, pre-trained models standardized based on the new dataset's mean and standard deviation outperformed those standardized with the pre-training dataset's statistics.

Training and validation loss curves revealed potential overfitting in U-Net and DeeplabV3 models, highlighted by fluctuations in validation loss. We propose integrating regularization techniques and data augmentation strategies to address this and improve generalization. Additionally, tuning hyperparameters and incorporating BatchNormalization layers to the U-Net could mitigate overfitting issues.

Our findings indicate that U-Net is a robust choice for wound segmentation tasks, especially when pre-trained and standardized appropriately. For mobile applications, LRASPP shows strong potential. Future work should refine these models for mobile deployment, ensuring reliable and effective outcomes in medical applications.

Keywords: image analysis, deep learning, convolutional neural networks, semantic segmentation, wound segmentation, medical image analysis, U-Net, DeeplabV3, LRASPP, PyTorch.

Problem statement

Our novel research on human skin wound segmentation is of significant importance in the medical field. Accurate segmentation is not just a task, but a crucial step in tracking wound healing status, aiding in treatment planning and monitoring. In skin grafting, segmentation is a key factor in assessing the wound area and condition, leading to improved surgical outcomes. Precise wound segmentation also allows healthcare providers to quantify and document wound characteristics, enhancing patient records and supporting research and clinical trials.

While image classification is a well-known task of predicting a class to which some specific image belongs, image segmentation is more complicated. It can be formulated as a classification problem of pixels with semantic labels (semantic segmentation) or partitioning of individual objects (instance segmentation) [1].

For a long time, semantic segmentation has been mainly performed using traditional methods like k-means clustering, watershed, etc. However, recent years have witnessed a surge of research, summarizing the rapid development of deep learning in the computer vision field and semantic segmentation in particular [1, 2]. Many studies reveal and develop the use of semantic segmentation in the medical field [3–6]. These advancements in deep learning techniques, particularly in the context of semantic segmentation, have shown remarkable promise in the medical field. They significantly contribute to computer-aided diagnosis systems and Convolutional Neural Networks (CNNs), providing high accuracy and reliability in wound segmentation tasks based on extensive research. This exciting progress opens up new possibilities for researchers in the field of medical image analysis.

Analysis of recent publications

Much research concentrates on solving segmentation problems in different subject areas: lung segmentation [7], cardiac segmentation [6], brain tumor segmentation [8], and others; this paper concentrates on the problem of wound image semantic segmentation. We compare modern CNN architectures that are applied to this task. Our selection criteria for these architectures include their performance in semantic segmentation tasks, their suitability for mobile device deployment, and their availability of pre-trained models. Because of the sensibility of medical data, we decided to introduce one more selection criterion - the ability to use models on mobile devices. The model's availability on mobile platforms can reduce the risks of using n-tier client-server architecture (where $n > 1$) by avoiding requests on the remote server and making all the computations on the local mobile device.

CNNs are used in almost every modern research on image processing. In [9], the authors present an approach where a deep convolutional neural network (AlexNet [10]) is used as a feature extractor, which receives pixel patches as inputs and then passes the extracted features to the SVM classifier. While this approach could be used for segmentation, the main goal was to perform wound classification. Authors of [11] propose a system of wound analysis, which includes wound segmentation, using convolutional autoencoder architecture that is used both to extract features for further classification using SVM and to perform wound image segmentation by creating probability maps from the output. Authors have reached 47.3% of mean Intersection over Union metric (IoU) on the test set. The significant problem was the need for annotated data, as it is common when researchers do not disclose the data they used. This problem has been alleviated by [12], where authors analyzed different convolutional architectures for wound semantic segmentation and have opened access to their labeled database of 1109 foot ulcer images. They received a best Dice score of 90.47% on their dataset and 94.05% on the labeled Medetec Wound Dataset [13]. We used a combination of their dataset and the Medetec foot ulcer data, which they labeled for the Foot Ulcer Segmentation Challenge, as our primary choice for the experiments. We were hoping to use the other wound-related dataset [14], but it is currently unavailable because of some error while trying to access it.

It is not uncommon to utilize transfer learning, a machine learning technique in which models pre-trained on images from one domain are used as a starting point to train models to perform similar tasks on the other domain.

Despite the functional task of segmenting the wound area accurately, the other important non-functional requirements are computational resource use and data privacy.

The work aims to evaluate and compare a few popular convolutional neural network (CNN) architectures for the semantic segmentation of human skin wounds, focusing on models suitable for mobile device deployment. This includes establishing a comparison framework to evaluate the pre-trained vs. not pre-trained models with different input data normalization parameters.

Presentation of the main material

Semantic segmentation evaluation metrics

Different metrics are used to evaluate model performance in semantic segmentation. The first widely used one is Intersection over Union (1), also called the Jaccard index.

$$IoU = \frac{TP}{TP+FP+FN}, \quad (1)$$

where TP, FP, and FN are the counts of true positive, false positive, and false negative pixels, respectively. For example, as we are dealing with the binary classification problem, true positive means that the positive (wound) class has been assigned to some pixel that really belongs to the positive (wound) class (so the prediction is true), according to the ground truth mask. The IOU metric can be interpreted as an intersection of predicted and ground truth labels divided by their union, which also follows from the metric name.

The other commonly used metrics are Precision (2) and Recall (3), which can be calculated as:

$$Precision = \frac{TP}{TP+FP}, \quad (2)$$

$$Recall = \frac{TP}{TP+FN}, \quad (3)$$

where TP, FP, and FN are calculated similarly for the IoU described above. These two metrics are interdependent; if we increase one, the second decreases, and vice versa. There are some tasks that need particularly false positives or false negatives to be as small as possible, and then one of the metrics is selected respectively. We decided to use the other popular metric in the field of image segmentation that represents the balanced approach between false positives and false negatives - Dice coefficient (4), or F1 score:

$$Dice = \frac{2TP}{2TP+FN+FP}, \quad (4)$$

which is a harmonic mean of Precision and Recall. The other sometimes used metric is pixel accuracy.

However, it can be very deceptive while having class imbalance problems, which is typical for segmentation tasks, where the background may occupy much of the image. Thus, we used the Dice score to evaluate our model performance and Dice loss as an error function. Using them, we can compare our results with those of the others.

Convolutional neural network architecture selection

There are multiple modern CNNs architecture. Authors of [15] state that the models they tested (FCN, ParseNet, Conv&Deconv, U-Net, FPN, Mask RCNN) require high memory and time consumption and that there is a problem with computation complexity. Thus, they must be more suitable for mobile devices with limited resources. Ways to solve this problem are described by [1], which proposes a few options:

- Use of simpler or designed explicitly for mobile device models.
- Use of model compression techniques.
- Use of knowledge distillation techniques.

The main idea of knowledge distillation is to learn a small student model from a large teacher model [16]. While this technique may give good results, its usage for wound segmentation tasks is separate research not covered in this paper.

Model compression is a beneficial set of techniques that can increase the resources needed for a model without significantly losing its quality. Two performant and popular techniques (some authors combine them into one category) are parameter pruning and quantization. According to [17], network quantization compresses the original network by reducing the number of bits required to represent each weight, and 8-bit and 16-bit quantization of the parameters can result in significant speed-up with minimal loss of accuracy. The main idea is to move from the floating-point weight representation to the integer one. Pruning is about reducing redundant parameters that are not sensitive to the performance [17]. This is effective while using pre-trained models

Backbone is a definition commonly used in segmentation models. Some segmentation networks build their structure from scratch, while others can reuse the left part (feature extractor) from the most suitable and efficient classification networks. That is called the backbone—the encoder part of the segmentation network that is a classification network itself without a final linear classifier probe layer. This allows us to select different backbones for the same segmentation decoder architecture to find the optimal one.

While implementing simpler models is always necessary, we must remember that some of the large ones were initially used for biomedical image segmentation.

U-Net [18]: the critical feature of this network, as well as lots of other modern segmentation networks, is that it has no fully connected layers, so we are not limited to some specific input image dimensions. Also, the authors state that U-Net applies to various biomedical segmentation problems, while initially, it has been developed to perform the segmentation of neuronal structures in EM stacks [18]. The network consists of the encoder and decoder (left and right parts) or contractive and expansive paths, as they are called in the original paper. While the left part is an ordinary convolutional network, the right one performs up sampling and has a concatenation of feature maps of the corresponding size from the left part. U-Net can also be used with different backbones, but we used a pre-trained model for the brain segmentation task [19].

Other models that have been considered are:

DeepLabV3 [20] with Dilated MobileNetV3 Large Backbone [20]. Atrous (dilated) convolution and Atrous Spatial Pyramid Pooling (ASPP) were the main features of the previous versions of this network. In the third version, authors enhance their previous approaches with cascaded or parallel atrous convolutions to capture multi-scale context by adopting multiple atrous rates and augmenting ASPP.

Lite R-ASPP [21] (Lite Reduced Atrous Spatial Pyramid Pooling) with Dilated MobileNetV3 Large Backbone. It is an improved version of the R-ASPP - segmentation part for MobileNetV3. It uses the global average pooling in a fashion like the Squeeze-and-Excitation. Results show that using it with the MonelNetV3 takes less than one second to segment an image with a resolution of 512 × 1024.

Dilated MobileNetV3 Large Backbone [21]. MobileNet [22] models were the first performant deep CNN for classification that was small enough to work on mobile devices. They introduced two sets of hyperparameters to choose from, depending on the available resources. The models were mostly built using depth-wise separable convolutions developed in [23]. The third version adds squeeze and excitation layers [24] to improve ImageNet's accuracy [25]. The top-1 accuracy of the V3-Large 1.0 version is 4.4% better than V2.1.0, with only 219 operations in comparison with 300 of the previous versions. The difference in parameters is 5.4M versus 3.4 in the V2.1.0.

Experiment setup

Experiments were carried out using the PyTorch and PyTorch Lightning library, which enhances PyTorch by providing a structured framework to conduct experiments and gather results efficiently. Training, validation, and 'workstation' testing of models were performed on Google Colaboratory PRO using a T4 GPU with 15GB of video memory.

The image segmentation pipeline consists of the following steps:

- Train/validation/test split
- Model training, validation and testing

Images from the Foot Ulcer Segmentation Challenge dataset [11] were divided into buckets of 503/200/107 images for the train/validation/test respectively.

Common conditions:

- 150 epochs of training

- Adam optimizer with an initial learning rate set to 0.001
- CosineAnnealingWarmRestarts learning rate scheduler with restarts every ten epochs and a minimal learning rate of 1e-7
- Batch size of 16
- Bo additional image augmentation
- Dice loss
- Gradients of pre-trained models are not frozen, allowing model weights to be adjusted during training

Table 1 summarizes the different experiment setups, detailing the use of pretrained versus not pretrained weights and the data standardization methods applied.

Table 1

Experiment setups details

Setup	Weights	Data Standardization
1	Pre-trained	Mean = [0.4079, 0.3448, 0.3187], std = [0.3005, 0.2517, 0.2409]
2	Not pre-trained	Mean = [0.4079, 0.3448, 0.3187], std = [0.3005, 0.2517, 0.2409]
3	Pre-trained	U-Net: mean = [0.0076, 0.0405, 0.0596], std = [0.0288, 0.0857, 0.0640] Other models: mean = [0.485, 0.456, 0.406], std = [0.229, 0.224, 0.225]
4	Pre-trained	No input data standardization
5	Not pre-trained	No input data standardization

- Setup 1: Uses pre-trained weights for all models with datasets standardized according to the mean and standard deviation of the Foot Ulcer Segmentation Challenge dataset.
- Setup 2: Uses, not pre-trained, weights for all models with the same dataset standardization as Setup 1.
- Setup 3: Uses pre-trained weights with datasets standardized according to the mean and standard deviation of the datasets on which the models were trained:
  - U-Net: Mean = [0.0076, 0.0405, 0.0596], Std = [0.0288, 0.0857, 0.0640]
  - Other Models: Mean = [0.485, 0.456, 0.406], Std = [0.229, 0.224, 0.225]
- Setup 4: Uses pre-trained weights with no input data standardization.
- Setup 5: Uses, not pre-trained weights with no input data standardization.

Results

The results of our experiments are analyzed based on the training and validation loss changes over epochs (Fig. 1) and the test Dice scores for different models and experiment setups (Fig. 2). Additionally, we have recorded the best result epochs for each model and setup (Table 2).

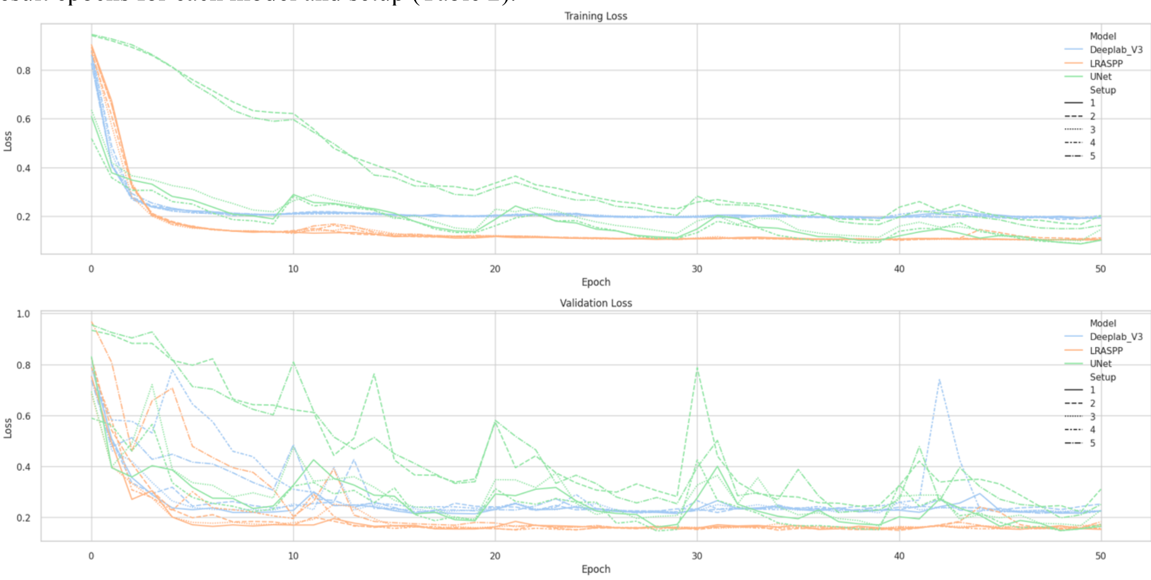

Fig. 1. Models train (top) and validation (bottom) loss change over time (epochs)

The training and validation loss plots (Fig. 1) indicate the following observations:

Training Loss:

- All models show a rapid decrease in training loss during the initial epochs, with the loss stabilizing around 20-30 epochs.
- The LRASPP models generally exhibit the lowest training loss across all setups, indicating good convergence and fitting to the training data.
- The U-Net models display moderate training loss, with some fluctuations, particularly in setups without standardization (Setups 4 and 5).
- The DeeplabV3 models have higher training loss compared to LRASPP and U-Net, suggesting potential challenges in fitting the training data.

Validation Loss:
- The validation loss shows more fluctuations compared to the training loss, which is expected due to the variations in the validation dataset.
- LRASPP models maintain relatively lower validation loss across all setups, similar to the training loss trend.
- U-Net models exhibit noticeable fluctuations in validation loss, indicating potential overfitting in certain setups.
- DeeplabV3 models also show fluctuations in validation loss, with some setups having higher peaks, suggesting instability and potential overfitting.

The bar plot (Fig. 2) presents the test Dice scores for different models and experiment setups:

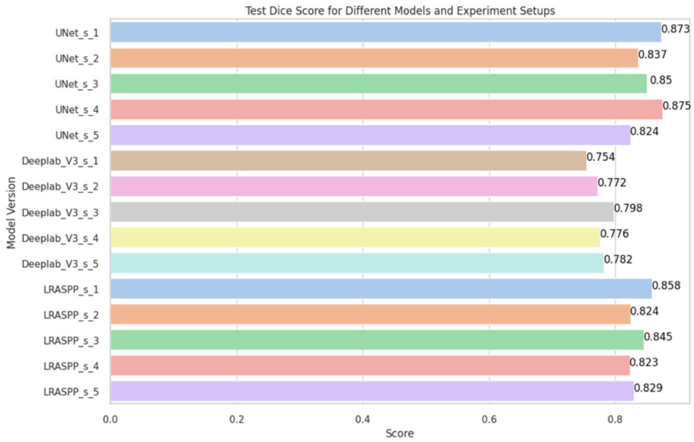

Fig. 2. Model test Dice score

Dice Score:
- The U-Net models achieve the highest Dice scores across most setups, with UNet_s_4 (0.875) being the best.
- LRASPP models also perform well, with LRASPP_s_1 (0.858) being the top performer among LRASPP setups.
- DeeplabV3 models have lower Dice scores compared to U-Net and LRASPP, with Deeplab_V3_s_3 (0.798) being the best among its setups.

The best result epochs for each model and setup are summarized in Table 2. This indicates the epoch at which each model achieved its highest Dice score during testing:

Table 2

Best Result Epochs

Model	Setup	Best Result Epoch
Deeplab_V3	**1**	19
Deeplab_V3	2	31
Deeplab_V3	3	39
Deeplab_V3	4	49
Deeplab_V3	5	67
LRASPP	1	148
LRASPP	2	52
LRASPP	3	92
LRASPP	4	67
LRASPP	5	105
UNet	1	111
UNet	2	145
UNet	3	124
UNet	4	119
UNet	5	144

Model Performance:
- U-Net models consistently achieve high Dice scores, indicating strong performance in wound segmentation tasks. This is evident in the highest Dice score of 0.875 for UNet_s_4.
- LRASPP models also perform well, particularly in pre-trained setups with data standardization, with LRASPP_s_1 achieving a Dice score of 0.858.
- DeeplabV3 models, while capable, lag behind U-Net and LRASPP in terms of segmentation accuracy, with the best performance from Deeplab_V3_s_3 at 0.798.

Effect of Pre-training and Standardization:
- Pre-trained models with appropriate data standardization (Setups 1 and 3) generally perform better across

all models. This can be seen in the high Dice scores for both U-Net and LRASPP models in these setups.
- Standardizing based on the new (Foot Ulcer dataset) dataset's mean and standard deviation (Setup 1) generally provides better performance than using the pretrained dataset's mean and standard deviation (Setup 3). This is particularly evident for U-Net, which achieves one of its best performances in Setup 1.
- Lack of data standardization (Setups 4 and 5) results in higher variability and potential overfitting, particularly in U-Net and DeeplabV3 models. For instance, UNet_s_5 and LRASPP_s_5 show lower performance compared to their standardized counterparts.
- U-Net has greater benefit from the use of pre-training weights, while having large fluctuations

Epochs and Overfitting:
- The best results are often achieved around 50-100 epochs, with some models benefiting (mostly U-Net) from longer training (up to 150 epochs). This is reflected in the best result epochs, where most models reach their peak performance within the first 100 epochs.
- Fluctuations in validation loss suggest potential overfitting, highlighting the need for regularization techniques and possibly more data augmentation.

Future steps to improve model performance:
- Regularization techniques such as dropout and weight decay, along with data augmentation strategies and learnable and adaptable data standardization techniques, like BatchNormalization layers for U-Net model, could help mitigate overfitting and enhance generalization.
- Expanding the dataset with more annotated images is essential to enhance model training and validation, leading to better segmentation accuracy.
- Increasing the dataset with more annotated images will enhance model training and validation accuracy.

Conclusions

This study successfully fulfills the aim of evaluating and comparing convolutional neural network (CNN) architectures suitable for mobile device deployment for the semantic segmentation of human skin wounds. Our analysis of the literature and experimental results highlights several key findings:

Model Comparison: We compared classical medical network U-Net with mobile-ready networks LRASPP and DeeplabV3. The U-Net models consistently achieved the highest Dice scores, indicating superior performance in wound segmentation tasks. LRASPP models also performed well, especially in pre-trained setups with data standardization, while DeeplabV3 showed lower segmentation accuracy.

Pre-training vs. Not Pre-training: Pre-trained models with appropriate data standardization generally outperformed those without pre-training. Standardizing based on the new dataset's mean and standard deviation provided better performance, especially for U-Net.

Data Standardization: Models that used data standardization showed better performance and stability compared to those without standardization. This highlights the importance of proper data preprocessing for achieving high segmentation accuracy.

Overfitting and Regularization: The observed fluctuations in validation loss suggest potential overfitting. We proposed integrating BatchNormalization layers in U-Net, using regularization techniques such as dropout and weight decay, and implementing data augmentation strategies to improve generalization.

Focusing on these aspects can help improve the reliability and effectiveness of wound segmentation models in medical applications.

References

1. Image segmentation using deep learning: a survey [Electronic resource] / Shervin Minaee [et al.] // IEEE transactions on pattern analysis and machine intelligence. – 2021. – P. 1. – Mode of access: https://doi.org/10.1109/tpami.2021.3059968 (date of access: 22.06.2024). – Title from screen.
2. RTSeg: real-time semantic segmentation comparative study [Electronic resource] / Mennatullah Siam [et al.] // 2018 25th IEEE international conference on image processing (ICIP), Athens, 7–10 October 2018. – [S. l.], 2018. – Mode of access: https://doi.org/10.1109/icip.2018.8451495 (date of access: 22.06.2024). – Title from screen.
3. Deep learning techniques for medical image segmentation: achievements and challenges [Electronic resource] / Mohammad Hesam Hesamian [et al.] // Journal of digital imaging. – 2019. – Vol. 32, no. 4. – P. 582–596. – Mode of access: https://doi.org/10.1007/s10278-019-00227-x (date of access: 22.06.2024). – Title from screen.
4. Interactive medical image segmentation using deep learning with image-specific fine tuning [Electronic resource] / Guotai Wang [et al.] // IEEE transactions on medical imaging. – 2018. – Vol. 37, no. 7. – P. 1562–1573. – Mode of access: https://doi.org/10.1109/tmi.2018.2791721 (date of access: 22.06.2024). – Title from screen.
5. Lai M. Deep learning for medical image segmentation [Electronic resource] / Matthew Lai // arXiv.org. – Mode of access: https://arxiv.org/abs/1505.02000 (date of access: 22.06.2024). – Title from screen.
6. Deep learning for cardiac image segmentation: a review [Electronic resource] / Chen Chen [et al.] // Frontiers in cardiovascular medicine. – 2020. – Vol. 7. – Mode of access: https://doi.org/10.3389/fcvm.2020.00025 (date of access: 22.06.2024). – Title from screen.
7. Deep learning with lung segmentation and bone shadow exclusion techniques for chest x-ray analysis of

lung cancer / Yuriy Gordienko [et al.] // Advances in computer science for engineering and education. – [S. l.], 2018. – P. 638–647.

8. Işın A. Review of mri-based brain tumor image segmentation using deep learning methods [Electronic resource] / Ali Işın, Cem Direkoğlu, Melike Şah // Procedia computer science. – 2016. – Vol. 102. – P. 317–324. – Mode of access: https://doi.org/10.1016/j.procs.2016.09.407 (date of access: 22.06.2024). – Title from screen.

9. Fine-Grained wound tissue analysis using deep neural network [Electronic resource] / H. Nejati [et al.] // ICASSP 2018 - 2018 IEEE international conference on acoustics, speech and signal processing (ICASSP), Calgary, AB, 15–20 April 2018. – [S. l.], 2018. – Mode of access: https://doi.org/10.1109/icassp.2018.8461927 (date of access: 22.06.2024). – Title from screen.

10. Alex Krizhevsky A. ImageNet classification with deep convolutional neural networks [Electronic resource] / Alex Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton // Communications of the ACM. – [S. l.], 2017. – P. 84–90. – Mode of access: https://doi.org/10.1145/3065386 (date of access: 22.06.2024). – Title from screen.

11. A unified framework for automatic wound segmentation and analysis with deep convolutional neural networks [Electronic resource] / Changhan Wang [et al.] // 2015 37th annual international conference of the IEEE engineering in medicine and biology society (EMBC), Milan, 25–29 August 2015. – [S. l.], 2015. – Mode of access: https://doi.org/10.1109/embc.2015.7318881 (date of access: 22.06.2024). – Title from screen.

12. Fully automatic wound segmentation with deep convolutional neural networks [Electronic resource] / Chuanbo Wang [et al.] // Scientific reports. – 2020. – Vol. 10, no. 1. – Mode of access: https://doi.org/10.1038/s41598-020-78799-w (date of access: 22.06.2024). – Title from screen.

13. Thomas S. Pictures of wounds and surgical wound dressings [Electronic resource] / Steve Thomas // Medetec Surgical Dressings and Wound Management Resource Centre - Home page. – Mode of access: http://www.medetec.co.uk/files/medetec-image-databases.html (date of access: 22.06.2024). – Title from screen.

14. Chronic wounds multimodal image database [Electronic resource] / Michał Kręcichwost [et al.] // Computerized medical imaging and graphics. – 2021. – Vol. 88. – P. 101844. – Mode of access: https://doi.org/10.1016/j.compmedimag.2020.101844 (date of access: 22.06.2024). – Title from screen.

15. Tapasvi B. A survey on semantic segmentation using deep learning techniques / B. Tapasvi, N. Udaya Kumar, E. Gnanamanoharan // Nternational journal of engineering research & technology (ijert). – Vol. 9, no. 5.

16. Knowledge distillation: a survey [Electronic resource] / Jianping Gou [et al.] // International journal of computer vision. – 2021. – Vol. 129, no. 6. – P. 1789–1819. – Mode of access: https://doi.org/10.1007/s11263-021-01453-z (date of access: 22.06.2024). – Title from screen.

17. A survey of model compression and acceleration for deep neural networks [Electronic resource] / Yu Cheng [et al.] // arXiv.org. – Mode of access: https://arxiv.org/abs/1710.09282 (date of access: 22.06.2024). – Title from screen.

18. Ronneberger O. U-Net: convolutional networks for biomedical image segmentation [Electronic resource] / O. Ronneberger, P. Fischer, T. Brox // Medical image computing and computer-assisted intervention – MICCAI 2015. MICCAI 2015. – [S. l.]. – Mode of access: https://doi.org/10.1007/978-3-319-24574-4_28 (date of access: 22.06.2024). – Title from screen.

19. Buda M. Association of genomic subtypes of lower-grade gliomas with shape features automatically extracted by a deep learning algorithm [Electronic resource] / Mateusz Buda, Ashirbani Saha, Maciej A. Mazurowski // Computers in biology and medicine. – 2019. – Vol. 109. – P. 218–225. – Mode of access: https://doi.org/10.1016/j.compbiomed.2019.05.002 (date of access: 22.06.2024). – Title from screen.

20. Rethinking atrous convolution for semantic image segmentation [Electronic resource] / Liang-Chieh Chen [et al.] // arXiv.org. – Mode of access: https://arxiv.org/abs/1706.05587 (date of access: 22.06.2024). – Title from screen.

21. Searching for MobileNetV3 [Electronic resource] / Andrew Howard [et al.] // arXiv.org. – Mode of access: https://arxiv.org/abs/1905.02244 (date of access: 22.06.2024). – Title from screen.

22. MobileNets: efficient convolutional neural networks for mobile vision applications [Electronic resource] / Andrew G. Howard [et al.] // arXiv.org. – Mode of access: https://arxiv.org/abs/1704.04861 (date of access: 22.06.2024). – Title from screen.

23. Sifre L. Rigid-Motion scattering for image classification [Electronic resource] : doctoral dissertation / Sifre Laurent. – Palaiseau, 2014. – 126 p. – Mode of access: https://www.di.ens.fr/data/publications/papers/phd_sifre.pdf (date of access: 22.06.2024). – Title from screen.

24. Hu J. Squeeze-and-Excitation networks [Electronic resource] / Jie Hu, Li Shen, Gang Sun // 2018 IEEE/CVF conference on computer vision and pattern recognition (CVPR), Salt Lake City, UT, 18–23 June 2018. – [S. l.], 2018. – Mode of access: https://doi.org/10.1109/cvpr.2018.00745 (date of access: 22.06.2024). – Title from screen.

25. ImageNet large scale visual recognition challenge [Electronic resource] / Olga Russakovsky [et al.] // International journal of computer vision. – 2015. – Vol. 115, no. 3. – P. 211–252. – Mode of access: https://doi.org/10.1007/s11263-015-0816-y (date of access: 22.06.2024). – Title from screen.