MIScnn: a framework for medical image segmentation with convolutional neural networks and deep learning

Dominik Müller, Frank Kramer

MIScnn: a framework for medical image segmentation with convolutional neural networks and deep learning

Dominik Müller[*] and Frank Kramer

Abstract

Background: The increased availability and usage of modern medical imaging induced a strong need for automatic medical image segmentation. Still, current image segmentation platforms do not provide the required functionalities for plain setup of medical image segmentation pipelines. Already implemented pipelines are commonly standalone software, optimized on a specific public data set. Therefore, this paper introduces the open-source Python library MIScnn.

Implementation: The aim of MIScnn is to provide an intuitive API allowing fast building of medical image segmentation pipelines including data I/O, preprocessing, data augmentation, patch-wise analysis, metrics, a library with state-of-the-art deep learning models and model utilization like training, prediction, as well as fully automatic evaluation (e.g. cross-validation). Similarly, high configurability and multiple open interfaces allow full pipeline customization.

Results: Running a cross-validation with MIScnn on the Kidney Tumor Segmentation Challenge 2019 data set (multi-class semantic segmentation with 300 CT scans) resulted into a powerful predictor based on the standard 3D U-Net model.

Conclusions: With this experiment, we could show that the MIScnn framework enables researchers to rapidly set up a complete medical image segmentation pipeline by using just a few lines of code. The source code for MIScnn is available in the Git repository: https://github.com/frankkramer-lab/MIScnn.

Keywords: Medical image analysis, Computer aided diagnosis, Biomedical image segmentation, U-Net, Deep learning, Open-source framework

Background

Medical imaging became a standard in diagnosis and medical intervention for the visual representation of the functionality of organs and tissues. Through the increased availability and usage of modern medical imaging like Magnetic Resonance Imaging (MRI) or Computed Tomography (CT), the need for automated processing of scanned imaging data is quite strong [1].

Currently, the evaluation of medical images is a manual process performed by physicians. Larger numbers of slices require the inspection of even more image material by doctors, especially regarding the increased usage of high-resolution medical imaging. In order to shorten the time-consuming inspection and evaluation process, an automatic pre-segmentation of abnormal features in medical images would be required.

Image segmentation is a popular sub-field of image processing within computer science [2–7]. The aim of semantic segmentation is to identify common features in an input image by learning and then labeling each pixel in

*Correspondence: dominik.mueller@informatik.uni-augsburg.de
IT-Infrastructure for Translational Medical Research, Faculty of Applied Computer Science, Faculty of Medicine, University of Augsburg, Augsburg, Germany

an image with a class (e.g. background, kidney or tumor). There is a wide range of algorithms to solve segmentation problems. However, state-of-the-art accuracy was accomplished by convolutional neural networks and deep learning models [7–11], which are used extensively today. Furthermore, the newest convolutional neural networks are able to exploit local and global features in images [12–14] and they can be trained to use 3D image information as well [15, 16]. In recent years, medical image segmentation models with a convolutional neural network architecture have become quite powerful and achieved similar results performance-wise as radiologists [10, 17]. Nevertheless, these models have been standalone applications with optimized architectures, preprocessing procedures, data augmentations and metrics specific for their data set and corresponding segmentation problem [14]. Also, the performance of such optimized pipelines varies drastically between different medical conditions. However, even for the same medical condition, evaluation and comparisons of these models are a persistent challenge due to the variety of the size, shape, localization and distinctness of different data sets. In order to objectively compare two segmentation model architectures from the sea of one-use standalone pipelines, each specific for a single public data set, it would be required to implement a complete custom pipeline with preprocessing, data augmentation and batch creation. Frameworks for general image segmentation pipeline building can not be fully utilized. The reason for this are their missing medical image I/O interfaces, their preprocessing methods, as well as their lack of handling highly unbalanced class distributions, which is standard in medical imaging. Recently developed medical image segmentation platforms, like NiftyNet [18], are powerful tools and an excellent first step for standardized medical image segmentation pipelines. However, they are designed more like configurable software instead of frameworks. They lack modular pipeline blocks to offer researchers the opportunity for easy customization and to help developing their own software for their specific segmentation problems.

In this work, we push towards constructing an intuitive and easy-to-use framework for fast setup of state-of-the-art convolutional neural network and deep learning models for medical image segmentation. The aim of our framework Medical Image Segmentation with Convolutional Neural Networks (MIScnn) is to provide a complete pipeline for preprocessing, data augmentation, patch slicing and batch creation steps in order to start straightforward with training and predicting on diverse medical imaging data. Instead of being fixated on one model architecture, MIScnn allows not only fast switching between multiple modern convolutional neural network models, but it also provides the possibility to easily add custom model architectures. Additionally, it facilitates a simple deployment and fast usage of new deep learning models for medical image segmentation. Still, MIScnn is highly configurable to adjust hyperparameters, general training parameters, preprocessing procedures, as well as include or exclude data augmentations and evaluation techniques.

Implementation
The open-source Python library MIScnn is a framework to setup medical image segmentation pipelines with convolutional neural networks and deep learning models. MIScnn is providing several core features, which are also illustrated in Fig. 1:

- 2D/3D medical image segmentation for binary and multi-class problems
- Data I/O, preprocessing and data augmentation for biomedical images
- Patch-wise and full image analysis
- State-of-the-art deep learning model and metric library
- Intuitive and fast model utilization (training, prediction)
- Multiple automatic evaluation techniques (e.g. cross-validation)
- Custom model, data I/O, pre-/postprocessing and metric support

Data input
NIfTI data I/O interface
MIScnn provides a data I/O interface for the Neuroimaging Informatics Technology Initiative (NifTI) [19] file format for loading Magnetic Resonance Imaging and Computed Tomography data into the framework. This format was initially created to speed up the development and enhance the utility of informatics tools related to neuroimaging. Still, it is now commonly used for sharing public and anonymous MRI and CT data sets, not only for brain imaging, but also for all kinds of human 3D imaging. A NIfTI file contains the 3D image matrix and diverse metadata, like the thickness of the MRI slices.

Custom data I/O interface
Next to the implemented NIfTI I/O interface, MIScnn allows the usage of custom data I/O interfaces for other imaging data formats. This open interface enables MIScnn to handle specific biomedical imaging features (e.g. MRI slice thickness), and therefore it avoids losing these feature information by a format conversion requirement. A custom I/O interface must be committed to the preprocessing function and it has to return the medical
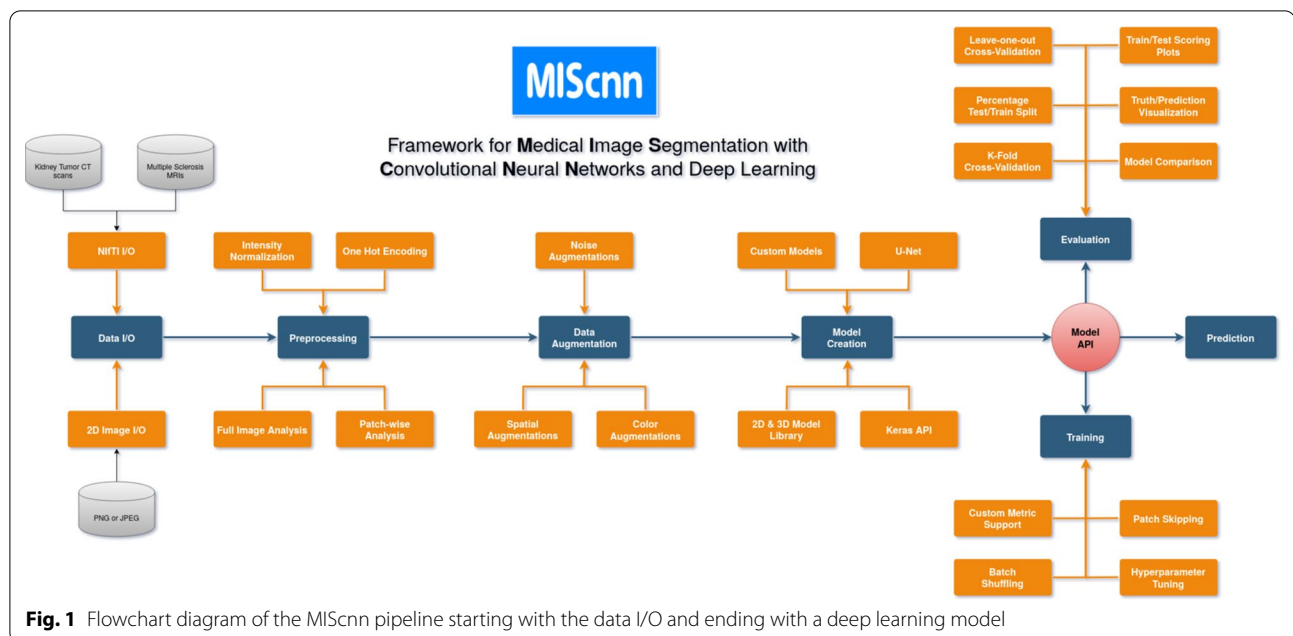
Fig. 1 Flowchart diagram of the MIScnn pipeline starting with the data I/O and ending with a deep learning model

image as a 2D or 3D matrix for integration in the work-flow. It is advised to add format specific preprocessing procedures (e.g. MRI slice thickness normalization) in the format specific I/O interface, before returning the image matrix into the pipeline.

Preprocessing
Pixel intensity normalization
Inconsistent signal intensity ranges of images can drastically influence the performance of segmentation methods [20, 21]. The signal ranges of biomedical imaging data are highly varying between data sets due to different image formats, diverse hardware/instruments (e.g. different scanners), technical discrepancies, and simply biological variation [10]. Additionally, the machine learning algorithms behind image segmentation usually perform better on features which follow a normal distribution. In order to achieve dynamic signal intensity range consistency, it is advisable to scale and standardize imaging data. The signal intensity scaling projects the original value range to a predefined range usually between [0, 1] or [−1, 1], whereas standardization centers the values close to a normal distribution by computing a Z-Score normalization. MIScnn can be configured to include or exclude pixel intensity scaling or standardization on the medical imaging data in the pipeline.

Clipping
Similar to pixel intensity normalization, it is also common to clip pixel intensities to a certain range. Intensity values outside of this range will be clipped to the minimum or

maximum range value. Especially in computer tomography images, pixel intensity values are expected to be identical for the same organs or tissue types even in different scanners [22]. This can be exploited through organ-specific pixel intensity clipping.

Resampling
The resampling technique is used to modify the width and/or height of images. This results into a new image with a modified number of pixels. Magnetic resonance or computer tomography scans can have different slice thickness. However, training neural network models requires the images to have the same slice thickness or voxel spacing. This can be accomplished through resampling. Additionally, downsampling images reduces the required GPU memory for training and prediction.

One hot encoding
MIScnn is able to handle binary (background/cancer) as well multi-class (background/kidney/liver/lungs) segmentation problems. The representation of a binary segmentation is being made quite simple by using a variable with two states, zero and one. But for the processing of multiple categorical segmentation labels in machine learning algorithms, like deep learning models, it is required to convert the classes into a more mathematical representation. This can be achieved with the One Hot encoding method by creating a single binary variable for

each segmentation class. MIScnn automatically One Hot encodes segmentation labels with more than two classes.

Patch-wise and full image analysis

Depending on the resolution of medical images, the available GPU hardware plays a large role in 3D segmentation analysis. Currently, it is not possible to fully fit high-resolution MRIs with an example size of $400 \times 512 \times 512$ into state-of-the-art convolutional neural network models due to the enormous GPU memory requirements. Therefore, the 3D medical imaging data can be either sliced into smaller cuboid patches or analyzed slice-by-slice, similar to a set of 2D images [10, 11, 23]. In order to fully use the information of all three axis, MIScnn slices 3D medical images into patches with a configurable size (e.g. $128 \times 128 \times 128$) by default. Depending on the model architecture, these patches can fit into GPUs with RAM sizes of 4–24 GB, which are commonly used in research. Nevertheless, the slice-by-slice 2D analysis, as well as the 3D patch analysis is supported and can be used in MIScnn. It is also possible to configure the usage of full 3D images in case of analyzing uncommonly small medical images or having a large GPU cluster. By default, 2D medical images are fitted completely into the convolutional neural network and deep learning models. Still, a 2D patch-wise approach for large resolution images can be also applied.

Data augmentation for training

In the machine learning field, data augmentation covers the artificially increase of training data. Especially in medical imaging, commonly only a small number of samples or images of a studied medical condition is available for training [10, 24–27]. Thus, an image can be modified with multiple techniques, like shifting, to expand the number of plausible examples for training. The aim is to create reasonable variations of the desired pattern in order to avoid overfitting in small data sets [26].

For state-of-the-art data augmentation, MIScnn integrated the batchgenerators package from the Division of Medical Image Computing at the German Cancer Research Center (DKFZ) [28]. It offers various data augmentation techniques and was used by the winners of the latest medical image processing challenges [14, 22, 29]. It supports spatial translations, rotations, scaling, elastic deformations, brightness, contrast, gamma and noise augmentations like Gaussian noise.

Sampling and batch generation
Skipping blank patches

The known problem in medical images of the large unbalance between the relevant segments and the background results into an extensive amount of parts purely labeled as background and without any learning information [10, 24]. Especially after data augmentation, there is no benefit to multiply these blank parts or patches [30]. Therefore, in the patch-wise model training, all patches, which are completely labeled as background, can be excluded in order to avoid wasting time on unnecessary fitting.

Batch management

After the data preprocessing and the optional data augmentation for training, sets of full images or patches are bundled into batches. One batch contains a number of prepared images which are processed in a single step by the model and GPU. Sequential for each batch or processing step, the neural network updates its internal weights accordingly with the predefined learning rate. The possible number of images inside a single batch highly depends on the available GPU memory and has to be configured properly in MIScnn. Every batch is saved to disk in order to allow fast repeated access during the training process. This approach drastically reduces the computing time due to the avoidance of unnecessary repeated preprocessing of the batches. Nevertheless, this approach is not ideal for extremely large data sets or for researchers without the required disk space. In order to bypass this problem, MIScnn also supports "on-the-fly" generation of the next batch in memory during runtime.

Batch shuffling

During model training, the order of batches, which are going to be fitted and processed, is shuffled at the end of each epoch. This method reduces the variance of the neural network during fitting over an epoch and lowers the risk of overfitting. Still, it must be noted, that only the processing sequence of the batches is shuffled and the data itself is not sorted into a new batch order.

Multi-CPU and -GPU support

MIScnn also supports the usage of multiple GPUs and parallel CPU batch loading next to the GPU computing. Particularly, the storage of already prepared batches on disk enables a fast and parallelizable processing with CPU as well as GPU clusters by eliminating the risk of batch preprocessing bottlenecks.

Deep learning model creation
Model architecture

The selection of a deep learning or convolutional neural network model is the most important step in a medical image segmentation pipeline. There is a variety of model architectures and each has different strengths and weaknesses [12, 13, 15, 16, 31–37]. MIScnn features an open model interface to load and switch between provided state-of-the-art convolutional neural network models like

the popular U-Net model [12]. Models are represented with the open-source neural network library Keras [38] which provides a user-friendly API for commonly used neural-network building blocks on top of TensorFlow [39]. The already implemented models are highly configurable by definable number of neurons, custom input sizes, optional dropout and batch normalization layers or enhanced architecture versions like the Optimized High Resolution Dense-U-Net model [15]. Additionally, MIScnn offers architectures for 3D, as well as 2D medical image segmentation. This model selection process is visualized in Fig. 2. Besides the flexibility in switching between already implemented models, the open model interface enables the ability for custom deep learning model implementations and simple integrating these custom models into the MIScnn pipeline.

Metrics

MIScnn offers a large quantity of various metrics which can be used as loss function for training or for evaluation in figures and manual performance analysis. The Dice coefficient, also known as the Dice similarity index, is one of the most popular metrics for medical image segmentation. It scores the similarity between the predicted segmentation and the ground truth. However, it also penalizes false positives comparable to the precision metric. Depending on the segmentation classes (binary or multi-class), there is a simple and class-wise Dice coefficient implementation. Whereas the simple implementation only accumulates the overall number of correct and false predictions, the class-wise implementation accounts the prediction performance for each segmentation class which is strongly recommended for commonly class-unbalanced medical images. Another popular supported metric is the Jaccard Index. Even though it is similar to the Dice coefficient, it does not only emphasize on precise segmentation. However, it also penalizes under- and over-segmentation. Still, MIScnn uses the Tversky loss [40] for training. Comparable to the Dice coefficient, the Tversky loss function addresses data imbalance. Even so, it achieves a much better trade-off between precision and recall. Thus, the Tversky loss function ensures good performance on binary, as well as multi-class segmentation. Additionally, all standard metrics which are included in Keras, like accuracy or cross-entropy, can be used in MIScnn. Next to the already implemented metrics or loss functions, MIScnn offers the integration of custom metrics for training and evaluation. A custom metric can be implemented as defined in Keras, and simply be passed to the deep learning model.

Model utilization

With the initialized deep learning model and the fully preprocessed data, the model can now be used for training on the data to fit model weights or for prediction by using an already fitted model. Alternatively, the model can perform an evaluation, as well, by running a cross-validation for example, with multiple training and prediction calls. The model API allows saving and loading models in order to subsequently reuse already fitted models for prediction or for sharing pre-trained models.

Training

In the process of training a convolutional neural network or deep learning model, diverse settings have to be configured. At this point in the pipeline, the data

Fig. 2 Flowchart visualization of the deep learning model creation and architecture selection process
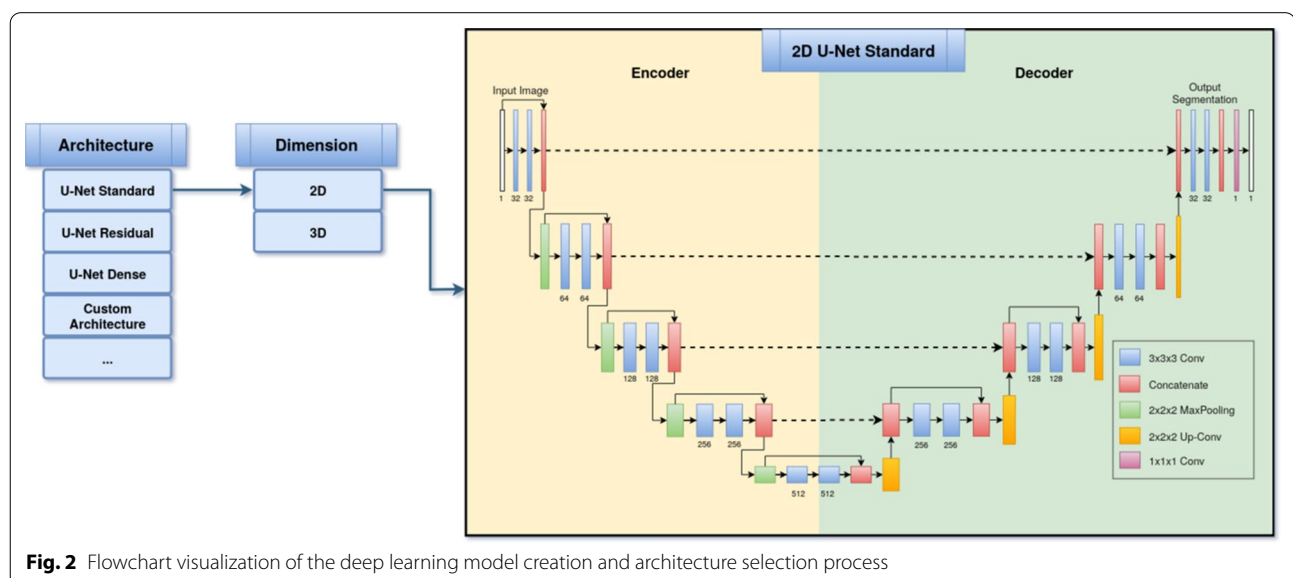
augmentation options of the data set, which have a large influence on the training in medical image segmentation, must be already defined. Sequentially, the batch management configuration covered the settings for the batch size, and also the batch shuffling at the end of each epoch. Therefore, only the learning rate and the number of epochs are required to be adjusted before running the training process. The learning rate of a neural network model is defined as the extend in which the old weights of the neural network model are updated in each iteration or epoch. In contrast, the number of epochs defines how many times the complete data set will be fitted into the model. Sequentially, the resulting fitted model can be saved to disk.

During the training, the underlying Keras framework gives insights into the current model performance with the predefined metrics, as well as the remaining fitting time. Additionally, MIScnn offers the usage of a fitting-evaluation callback functionality in which the fitting scores and metrics are stored into a tab separated file or directly plotted as a figure.

Prediction

For the segmentation prediction, an already fitted neural network model can be directly used after training or it can be loaded from file. The model predicts for every pixel a Sigmoid value for each class. The Sigmoid value represents a probability estimation of this pixel for the associated label. Sequentially, the argmax of the One Hot encoded class are identified for multi-class segmentation problems and then converted back to a single result variable containing the class with the highest Sigmoid value.

When using the overlapping patch-wise analysis approach during the training, MIScnn supports two methods for patches in the prediction. Either the prediction process plainly creates distinct patches and treats the overlapping patches during the training as purely data augmentation, or overlapping patches are created for prediction. Due to the lack of prediction power at patch edges, computing a second prediction for edge pixels in patches, by using an overlap, is a commonly used approach. In the following merge of patches back to the original medical image shape, a merging strategy for the pixels is required, in the overlapping part of two patches and with multiple predictions. By default, MIScnn calculates the mean between the predicted Sigmoid values for each class in every overlapping pixel.

The resulting image matrix with the segmentation prediction, which has the identical shape as the original medical image, is saved into a file structure according to the provided data I/O interface. By default, using the NIfTI data I/O interface, the predicted segmentation

matrix is saved in NIfTI format without any additional metadata.

Evaluation

MIScnn supports multiple automatic evaluation techniques to investigate medical image segmentation performance: k-fold cross-validation, leave-one-out cross-validation, percentage-split validation, hold-out sets for testing (data set split into test and train set with a given percentage) and detailed validation in which it can be specified which images should be used for training and testing. Except for the detailed validation, all other evaluation techniques use random sampling to create training and testing data sets. During the evaluation, the predefined metrics and loss function for the model are automatically plotted in figures and saved in tab separated files for possible further analysis. Next to the performance metrics, the pixel value range and segmentation class frequency of medical images can be analyzed in the MIScnn evaluation. Also, the resulting prediction can be compared directly next to the ground truth by creation image visualizations with segmentation overlays. For 3D images, like MRIs, the slices with the segmentation overlays are automatically visualized in the Graphics Interchange Format (GIF).

Results

Here, we analyze and evaluate data from the Kidney Tumor Segmentation Challenge 2019 using MIScnn. The main idea for this experiment is to demonstrate the 'out-of-the-box' performance of MIScnn without thorough and time-consuming optimization on the data set or on the medical abnormality. All results were obtained using the scripts shown in the Appendix.

Kidney Tumor Segmentation Challenge 2019 (KiTS19)

With more than 400,000 kidney cancer diagnoses worldwide in 2018, kidney cancer is under the top 10 most common cancer types in men and under the top 15 in woman [41]. The development of advanced tumor visualization techniques is highly important for efficient surgical planning. Due to the variety in kidney and kidney tumor morphology, the automatic image segmentation is challenging but of great interest [29].

The goal of the KiTS19 challenge is the development of reliable and unbiased kidney and kidney tumor semantic segmentation methods [29]. Therefore, the challenge built a data set for arterial phase abdominal CT scan of 300 kidney cancer patients [29]. The original scans have an image resolution of 512×512 and on average 216 slices (highest slice number is 1059). For all CT scans, a ground truth semantic segmentation was created by

Fig. 3 Computed Tomography image from the Kidney Tumor Segmentation Challenge 2019 data set showing the ground truth segmentation of both kidneys (red) and tumor (blue) [29]

experts. This semantic segmentation labeled each pixel with one of three classes: Background, kidney or tumor. An example CT scan including annotation is shown in Fig. 3. 210 of these CT scans with the ground truth segmentation were published during the training phase of the challenge, whereas 90 CT scans without published ground truth were released afterwards in the submission phase. The submitted user predictions for these 90 CT scans will be objectively evaluated and the user models ranked according to their performance. The CT scans were provided in NIfTI format in original resolution and also in interpolated resolution with slice thickness normalization.

Validation on the KiTS19 data set with MIScnn

For the evaluation of the MIScnn framework usability and data augmentation quality, a subset of 120 CT scans with slice thickness normalization were retrieved from the KiTS19 data set. An automatic threefold cross-validation was run on this KiTS19 subset with MIScnn. In order to reduce the overfitting risk, the cross-validation testing sets had no influence on the fitting process and were not used for any automatic hyper parameter tuning.

MIScnn configurations

The MIScnn pipeline was configured to perform a multi-class, patch-wise analysis with $80 \times 160 \times 160$ patches and a batch size of 2. The pixel value normalization by Z-Score, clipping to the range -79 and 304, as well as resampling to the voxel spacing $3.22 \times 1.62 \times 1.62$.

For data augmentation, all implemented techniques were used. This includes creating patches through random cropping, scaling, rotations, elastic deformations, mirroring, brightness, contrast, gamma and Gaussian noise augmentations. For prediction, overlapping patches were created with an overlap size of $40 \times 80 \times 80$ in x, y, z directions. The standard 3D U-Net with batch normalization layers were used as deep learning and convolutional neural network model. The training was performed using the Tversky loss for 1000 epochs with a starting learning rate of 1E−4 and batch shuffling after each epoch.

The cross-validation was run on two Nvidia Quadro P6000 (24 GB memory each), using 48 GB memory and taking 58 h.

Results

With the MIScnn pipeline, it was possible to successfully set up a complete, working medical image multi-class segmentation pipeline. The threefold cross-validation of 120 CT scans for kidney and tumor segmentation were evaluated through several metrics: Tversky loss, soft Dice coefficient, class-wise Dice coefficient, as well as the sum of categorical cross-entropy and soft Dice coefficient. These scores were computed during the fitting itself, as well as for the prediction with the fitted model. For each cross-validation fold, the training and predictions scores are visualized in Fig. 4 and sum up in Table 1.

The fitted model achieved a very strong performance for kidney segmentation. The kidney Dice coefficient had a median around 0.9544. The tumor segmentation prediction showed a considerably high but weaker performance than the kidney with a median around 0.7912.

Besides the computed metrics, MIScnn created segmentation visualizations for manual comparison between ground truth and prediction. As illustrated in Fig. 5, the predicted semantic segmentation of kidney and tumors is highly accurate.

Discussion
MIScnn framework

With excellent performing convolutional neural network and deep learning models like the U-Net, the urge to move automatic medical image segmentation from the research labs into practical application in clinics is uprising. Still, the landscape of standalone pipelines of top performing models, designed only for a single specific public data set, handicaps this progress. The goal of MIScnn is to provide a high-level API to setup a medical image segmentation pipeline with preprocessing, data augmentation, model architecture selection and model utilization. MIScnn offers a highly configurable and open-source pipeline with several interfaces for custom deep learning models, image formats or fitting metrics.

Fig. 4 Performance evaluation of the standard 3D U-Net model for kidney and tumor prediction with a threefold cross-validation on the 120 CT data set from the KiTS19 challenge. Left: Tversky loss against epochs illustrating loss development during training for the corresponding test and train data sets. Each point represents the average Tversky loss between the cross-validation folds. Center: Class-wise Dice coefficient against epochs illustrating soft Dice similarity coefficient development during training for the corresponding test and train data sets. Each point represents the average soft Dice similarity coefficient between the cross-validation folds. Right: Dice similarity coefficient distribution for the kidney and tumor for all samples of the cross-validation

Table 1 Performance results of the threefold cross-validation for tumor and kidney segmentation with a standard 3D U-Net model on 120 CT scans from the KiTS19 challenge
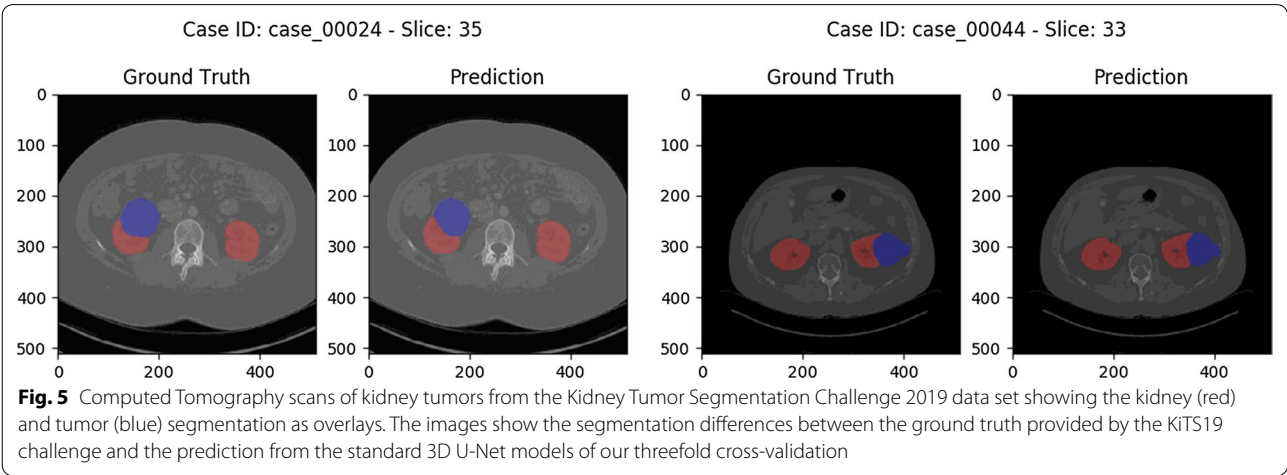
Metric	Training	Validation
Tversky loss	0.3672	0.4609
Soft Dice similarity coefficient	0.8776	0.8235
Categorical cross-entropy	− 0.8584	− 0.7899
Dice similarity coefficient: background	–	0.9994
Dice similarity coefficient: kidney	–	0.9319
Dice similarity coefficient: tumor	–	0.6750

Each metric is computed between the provided ground truth and our model prediction and then averaged between the three folds

The modular structure of MIScnn allows a medical image segmentation novice to setup a functional pipeline for a custom data set in just a few lines of code. Additionally, switchable models and an automatic evaluation functionality allow robust and unbiased comparisons between deep learning models. A universal framework for medical image segmentation, following the Python philosophy of simple and intuitive modules, is an important step in contributing to practical application development.

Use case: Kidney Tumor Segmentation Challenge

In order to show the reliability of MIScnn, a pipeline was setup for kidney tumors segmentation on a CT image

Fig. 5 Computed Tomography scans of kidney tumors from the Kidney Tumor Segmentation Challenge 2019 data set showing the kidney (red) and tumor (blue) segmentation as overlays. The images show the segmentation differences between the ground truth provided by the KiTS19 challenge and the prediction from the standard 3D U-Net models of our threefold cross-validation

data set. The popular and state-of-the-art standard U-Net were used as deep learning model with up-to-date data augmentation. Its predictive power was very impressive in the context of using only the standard U-Net architecture with mostly default hyperparameters. In the medical perspective, through the variety in kidney tumor morphology, which is one of the reasons for the KiTS19 challenge, the weaker tumor results are quite reasonable [29]. Also, the models were trained with only 38% of the original KiTS19 data set due to 80 images for training and 40 for testing were randomly selected. The remaining 90 CTs were excluded in order to reduce run time in the cross-validation. Nevertheless, it was possible to build a powerful pipeline for kidney tumor segmentation with MIScnn resulting into a model with high performance, which is directly comparable with modern, optimized, standalone pipelines [12, 13, 16, 32]. We proved that with just a few lines of codes using the MIScnn framework, it was possible to successfully build a powerful pipeline for medical image segmentation. Additionally, fast switching the model to a more precise architecture for high resolution images, like the Dense U-Net model, would probably result into an even better performance [15]. However, this gain would go hand in hand with an increased fitting time and higher GPU memory requirement, which was not possible with our current sharing schedule for GPU hardware. Nevertheless, the possibility of swift switching between models to compare their performance on a data set is a promising step forward in the field of medical image segmentation.

Road map and future direction

The active MIScnn development is currently focused on multiple key features: Adding further data I/O interfaces for the most common medical image formats like DICOM, extend preprocessing and data augmentation methods, implement more efficient patch skipping techniques instead of excluding every blank patch (e.g. denoising patch skipping) and implementation of an open interface for custom preprocessing techniques for specific image types like MRIs. Next to the planned feature implementations, the MIScnn road map includes the model library extension with more state-of-the-art deep learning models for medical image segmentation. Additionally, an objective comparison of the U-Net model version variety is outlined to get more insights on different model performances with the same pipeline. Community contributions in terms of implementations or critique are welcomed and can be included after evaluation. Currently, MIScnn already offers a robust pipeline for medical image segmentation, nonetheless, it will still be regularly updated and extended in the future.

MIScnn availability

The MIScnn framework can be directly installed as a Python library using pip install miscnn. Additionally, the source code is available in the Git repository: https ://github.com/frankkramer-lab/MIScnn. MIScnn is licensed under the open-source GNU General Public License Version 3. The code of the cross-validation experiment for the Kidney Tumor Segmentation Challenge is available as a Jupyter Notebook in the official Git repository.

Conclusions

In this paper, we have introduced the open-source Python library MIScnn: A framework for medical image segmentation with convolutional neural networks and deep learning. The intuitive API allows fast building medical image segmentation pipelines including data I/O, preprocessing, data augmentation, patch-wise analysis, metrics, a library with state-of-the-art deep learning models and model utilization like training, prediction, as well as fully automatic evaluation (e.g. cross-validation). High configurability and multiple open interfaces allow users to fully customize the pipeline. This framework enables researchers to rapidly set up a complete medical image segmentation pipeline by using just a few lines of code. We proved the MIScnn functionality by running an automatic cross-validation on the Kidney Tumor Segmentation Challenge 2019 CT data set resulting into a powerful predictor. We hope that it will help migrating medical image segmentation from the research labs into practical applications.

Availability and requirements

Project name: MIScnn
Project home page: https://github.com/frankkramer-lab/MIScnn
Operating system(s): Platform independent
Programming language: Python
Other requirements: Python 3.6; Tensorflow 2.2.0
License: GPL-3.0 License
Any restrictions to use by non-academics: None

Abbreviations
MIScnn: Medical Image Segmentation with Convolutional Neural Networks; CT: Computed tomography; I/O: Input/output; MRI: Magnetic Resonance Imaging; NifTI: Neuroimaging Informatics Technology Initiative; GPU: Graphics processing unit; RAM: Random Access Memory; CPU: Central Processing Unit; API: Application programming interface; KiTS19: Kidney Tumor Segmentation Challenge 2019.

Availability of data and materials
The KiTS19 dataset analyzed in this article is available in the GitHub repository, https://github.com/neheller/kits19. All data generated and analyzed during this study is available in the zenodo repository, https://doi.org/10.5281/zenodo.3962097. The code for this article was implemented in Python (platform independent) and is available under the GPL-3.0 License at the following GitHub repository: https://github.com/frankkramer-lab/MIScnn.

Ethics approval and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
The authors declare no conflicts of interest.

Appendix
Please check the Appendix document in the below mentioned path

> Direct URL: https://zenodo.org/record/3962097
> DOI URL: https://doi.org/10.5281/zenodo.3962097

References
1. Aggarwal P, Vig R, Bhadoria S, Dethe CG. Role of segmentation in medical imaging: a comparative study. Int J Comput Appl. 2011;29:54–61.
2. Gibelli D, Cellina M, Gibelli S, Oliva AG, Termine G, Pucciarelli V, et al. Assessing symmetry of zygomatic bone through three-dimensional segmentation on computed tomography scan and "mirroring" procedure: a contribution for reconstructive maxillofacial surgery. J Cranio-Maxillofac Surg. 2018;46:600–4. https://doi.org/10.1016/j.jcms.2018.02.012.
3. Cellina M, Gibelli D, Cappella A, Toluian T, Pittino CV, Carlo M, et al. Segmentation procedures for the assessment of paranasal sinuses volumes. Neuroradiol J. 2020. https://doi.org/10.1177/1971400920946635.
4. Hu X, Luo W, Hu J, Guo S, Huang W, Scott MR, et al. Brain SegNet: 3D local refinement network for brain lesion segmentation. BMC Med Imaging. 2020;20:17. https://doi.org/10.1186/s12880-020-0409-2.
5. Sun R, Wang K, Guo L, Yang C, Chen J, Ti Y, et al. A potential field segmentation based method for tumor segmentation on multi-parametric MRI

6. of glioma cancer patients. BMC Med Imaging. 2019;19:48. https://doi.org/10.1186/s12880-019-0348-y.
6. Claudia C, Farida C, Guy G, Marie-Claude M, Carl-Eric A. Quantitative evaluation of an automatic segmentation method for 3D reconstruction of intervertebral scoliotic disks from MR images. BMC Med Imaging. 2012;12:26. https://doi.org/10.1186/1471-2342-12-26.
7. Guo Y, Liu Y, Georgiou T, Lew MS. A review of semantic segmentation using deep neural networks. Int J Multimed Inf Retr. 2018;7:87–93. https://doi.org/10.1007/s13735-017-0141-z.
8. Anwar SM, Majid M, Qayyum A, Awais M, Alnowami M, Khan MK. Medical image analysis using convolutional neural networks: a review. J Med Syst. 2018;42:226. https://doi.org/10.1007/s10916-018-1088-1.
9. Wang G. A perspective on deep imaging. IEEE Access. 2016;4:8914–24.
10. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. Med Image Anal. 2012;2017(42):60–88.
11. Shen D, Wu G, Suk H-I. Deep learning in medical image analysis. Annu Rev Biomed Eng. 2017;19:221–48. https://doi.org/10.1146/annurev-bioeng-071516-044442.
12. Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. Lect Notes Comput Sci (including SubserLect Notes ArtifIntellLect Notes Bioinformatics). 2015;9351:234–41.
13. Zhou Z, Siddiquee MMR, Tajbakhsh N, Liang J. UNet++: a nested U-Net architecture for medical image segmentation. 2018. http://arxiv.org/abs/1807.10165. Accessed 19 Jul 2019.
14. Isensee F, Petersen J, Klein A, Zimmerer D, Jaeger PF, Kohl S, et al. nnU-Net: self-adapting framework for U-Net-based medical image segmentation. 2018. http://arxiv.org/abs/1809.10486. Accessed 19 Jul 2019.
15. Kolařík M, Burget R, Uher V, Říha K, Dutta M. Optimized high resolution 3D dense-U-Net network for brain and spine segmentation. Appl Sci. 2019;9:404.
16. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O. 3D U-net: learning dense volumetric segmentation from sparse annotation. Lect Notes Comput Sci. 2016;9901:424–32.
17. Lee K, Zung J, Li P, Jain V, Seung HS. Superhuman accuracy on the SNEMI3D connectomics challenge. 2017; Nips:1–11. http://arxiv.org/abs/1706.00120.
18. Gibson E, Li W, Sudre C, Fidon L, Shakir DI, Wang G, et al. NiftyNet: a deep-learning platform for medical imaging. Comput Methods Programs Biomed. 2018;158:113–22. https://doi.org/10.1016/j.cmpb.2018.01.025.
19. Neuroimaging Informatics Technology Initiative. https://nifti.nimh.nih.gov/background. Accessed 19 Jul 2019.
20. Roy S, Carass A, Prince JL. Patch based intensity normalization of brain MR images. In: Proceedings—international symposium on biomedical imaging. 2013.
21. Nyú LG, Udupa JK. On standardizing the MR image intensity scale. Magn Reson Med. 1999;42:1072–81. https://doi.org/10.1002/(SICI)1522-2594(199912)42:6<1072::AID-MRM11>3.0.CO;2-M
22. Isensee F, Maier-Hein KH. An attempt at beating the 3D U-Net. 2019;1–7. http://arxiv.org/abs/1908.02182.
23. Lin G, Shen C, Van Den HA, Reid I. Efficient piecewise training of deep structured models for semantic segmentation. Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit. 2016;2016:3194–203.
24. Hussain Z, Gimenez F, Yi D, Rubin D. Differential data augmentation techniques for medical imaging classification tasks. Annu Symp Proc AMIA Symp. 2017;2017:979–84.
25. Eaton-rosen Z, Bragman F. Improving data augmentation for medical image segmentation. Midl. 2018; 1–3.
26. Perez L, Wang J. The effectiveness of data augmentation in image classification using deep learning. 2017. http://arxiv.org/abs/1712.04621. Accessed 23 Jul 2019.
27. Taylor L, Nitschke G. Improving deep learning using generic data augmentation. 2017. http://arxiv.org/abs/1708.06020. Accessed 23 Jul 2019.
28. Isensee F, Jäger P, Wasserthal J, Zimmerer D, Petersen J, Kohl S, et al. batchgenerators—a python framework for data augmentation. 2020. https://doi.org/10.5281/zenodo.3632567.
29. Heller N, Sathianathen N, Kalapara A, Walczak E, Moore K, Kaluzniak H, et al. The KiTS19 challenge data: 300 kidney tumor cases with clinical context, CT semantic segmentations, and surgical outcomes. 2019. http://arxiv.org/abs/1904.00445. Accessed 19 Jul 2019.

30. Coupé P, Manjón JV, Fonov V, Pruessner J, Robles M, Collins DL. Patch-based segmentation using expert priors: application to hippocampus and ventricle segmentation. Neuroimage. 2011;54:940–54. https://doi.org/10.1016/j.neuroimage.2010.09.018.

31. Wang G, Li W, Ourselin S, Vercauteren T. Automatic brain tumor segmentation using cascaded anisotropic convolutional neural networks. Lect Notes Comput Sci (including SubserLect Notes ArtifIntellLect Notes Bioinformatics). 2018;10670:178–90.

32. Zhang Z, Liu Q, Wang Y. Road extraction by deep residual U-Net. IEEE Geosci Remote Sens Lett. 2018;15:749–53.

33. Iglovikov V, Shvets A. TernausNet: U-Net with VGG11 encoder pre-trained on ImageNet for image segmentation. 2018. http://arxiv.org/abs/1801.05746. Accessed 19 Jul 2019.

34. Ibtehaz N, Rahman MS. MultiResUNet : Rethinking the U-Net architecture for multimodal biomedical image segmentation. 2019. http://arxiv.org/abs/1902.04049. Accessed 19 Jul 2019.

35. Kamnitsas K, Bai W, Ferrante E, McDonagh S, Sinclair M, Pawlowski N, et al. Ensembles of multiple models and architectures for robust brain tumour segmentation. Lect Notes Comput Sci (including SubserLect Notes ArtifIntellLect Notes Bioinformatics). 2018;10670:450–62.

36. Valverde S, Salem M, Cabezas M, Pareto D, Vilanova JC, Ramió-Torrentà L, et al. One-shot domain adaptation in multiple sclerosis lesion segmentation using convolutional neural networks. NeuroImage Clin. 2018;2019(21):101638. https://doi.org/10.1016/j.nicl.2018.101638.

37. Brosch T, Tang LYW, Yoo Y, Li DKB, Traboulsee A, Tam R. Deep 3D convolutional encoder networks with shortcuts for multiscale feature integration applied to multiple sclerosis lesion segmentation. IEEE Trans Med Imaging. 2016;35:1229–39.

38. Chollet, François, others. Keras. 2015. https://keras.io. Accessed 19 Jul 2019.

39. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, et al. Tensor-Flow: large-scale machine learning on heterogeneous systems. 2015. https://www.tensorflow.org/. Accessed 19 Jul 2019.

40. Seyed SSM, Erdogmus D, Gholipour A, Salehi SSM, Erdogmus D, Gholipour A. Tversky loss function for image segmentation using 3D fully convolutional deep networks. In: Lecture notes in computer science. Springer; 2017. p. 379–87. https://doi.org/10.1007/978-3-319-67389-9_44.

41. Ferlay J, Colombet M, Soerjomataram I, Mathers C, Parkin DM, Piñeros M, et al. Estimating the global cancer incidence and mortality in 2018: GLOBOCAN sources and methods. Int J Cancer. 2018. https://doi.org/10.1002/ijc.31937.

Publisher's Note