UNIVERSITY OF COPENHAGEN

FACULTY OF SCIENCE

One Network To Segment Them All

A General, Lightweight System for Accurate 3D Medical Image Segmentation

Mathias Perslev[1], Erik B. Dam[1,2], Akshay Pai[1,2] & Christian Igel[1]

[1] Department of Computer Science, University of Copenhagen, Denmark

[2] Cerebriu A/S, Copenhagen, Denmark

Abstract

Many recent medical segmentation systems rely on powerful deep learning models to solve highly specific tasks. To maximize performance, it is standard practice to evaluate numerous pipelines with varying model topologies, optimization parameters, pre- & postprocessing steps, and even model cascades. It is often not clear how the resulting pipeline transfers to different tasks.

We propose a simple and thoroughly evaluated deep learning framework for segmentation of arbitrary medical image volumes. The system requires no task-specific information, no human interaction and is based on a fixed model topology and a fixed hyperparameter set, eliminating the process of model selection and its inherent tendency to cause method-level over-fitting. The system is available in open source and does not require deep learning expertise to use. Without task-specific modifications, the system performed better than or similar to highly specialized deep learning methods across 3 separate segmentation tasks. In addition, it ranked 5-th and 6-th in the first and second round of the 2018 Medical Segmentation Decathlon comprising another 10 tasks.

The system relies on multi-planar data augmentation which facilitates the application of a single 2D architecture based on the familiar U-Net. Multi-planar training combines the parameter efficiency of a 2D fully convolutional neural network with a systematic train- and test-time augmentation scheme, which allows the 2D model to learn a representation of the 3D image volume that fosters generalization.

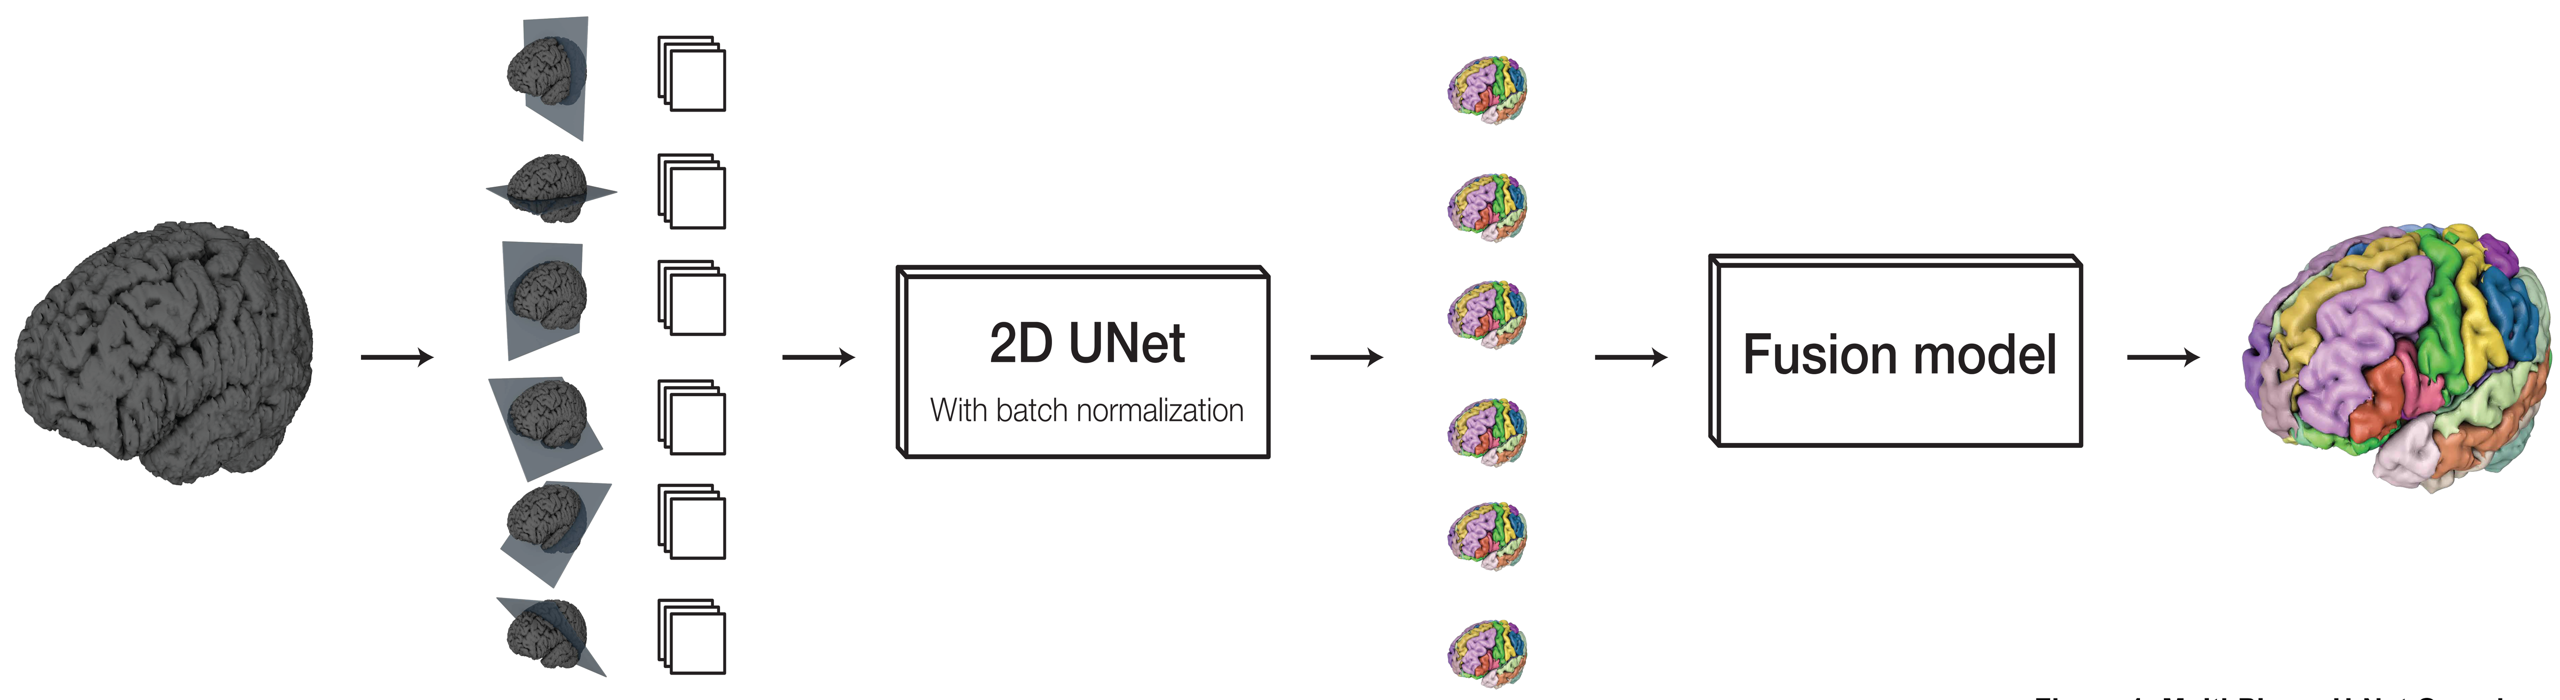

Figure 1: Multi-Planar U-Net Overview

Schematic overview of the Multi-Planar U-Net in the inference phase. The input volume (left) is sampled on 2D isotropic grids along multiple planar axes. The model predicts a full volume along each axis and maps the predictions into the original image space. A fusion model combines the 6 proposed segmentation volumes into a single final segmentation.

Method

- We propose the '*Multi-Planar U-Net*'; a general-purpose, lightweight system for 3D medical image segmentation.
- *Multi-planar augmentation:* A 2D fully convolutional network (e.g. U-Net[1]) is fit to isotropic images sampled across multiple planes.
- At test-time, the image volume is segmented in each plane. The segmentations are combined by a small *fusion model* that accounts for how easily each class is segmented in each plane.
- Multi-Planar U-Nets provide accurate segmentations across highly variable tasks *without hyperparameter tuning.*

Datasets & Evaluation

- We applied the Multi-Planar U-Net across 13 different medical image segmentation tasks.
- No modifications were made to the architecture or hyperparameters between tasks.
- 10 of the tasks were part of the 2018 Medical Segmentation Decathlon[2] challenge.
- We compared the Multi-Planar U-Net to a 3D U-Net trained with rotation augmentation as well as to an ensemble of single-plane 2D U-Nets.

Medical Segmentation Decathlon

http://medicaldecathlon.com/

- A rolling challenge that aims to evaluate the performance of general-purpose segmentation systems across tasks of significant variability.
- Participating systems must be autonomous, solving each task (separately) without human interaction.
- **Other noteworthy participants:** nnU-Net[3], K.A.V.athlon, NVDLMED, Lupin, LS Wang's Group, MIMI and others.

Results

- The Multi-Planar U-Net reached state-of-the-art performance for deep learning based methods on 3 different tasks (MICCAI[4], HarP[5] & OAI[6]).
- It ranked 5th and 6th place in the first and second phases of the 2018 Medical Segmentation Decathlon, comparing unfavorable mostly to systems that rely on compute intensive cross-validation experiments or model cascades.
- Multi-Planar U-Nets may be superior to 3D U-Nets and ensembles of single-plane 2D U-Nets.

	Dataset	Modality	Segmentation Target(s)	Classes	Size	F1 Score
	MICCAI	MRI	Whole-Brain	135	35	0.74 ± 0.03
	HarP	MRI	L+R Hippocampus	3	135	0.85 ± 0.03
	OAI	MRI	Knee Cartilages	7	176	0.87 ± 0.06
2018 Medical Segmentation Decathlon	Task 1	MRI	Brain Tumours	4	750	0.60 ± 0.24
	Task 2	MRI	Cardiac, Left Atrium	2	30	0.89 ± 0.09
	Task 3	CT	Liver & Tumour	2	201	0.76 ± 0.18
	Task 4	MRI	Hippocampus ROI.	2	394	0.89 ± 0.04
	Task 5	MRI	Prostate	3	48	0.78 ± 0.10
	Task 6	CT	Lung Tumours	2	96	0.59 ± 0.23
	Task 7	CT	Pancreas & Tumour	3	420	0.48 ± 0.21
	Task 8	CT	Hepatic Ves. & Tumour	3	443	0.49
	Task 9	CT	Spleen	2	61	0.95
	Task 10	CT	Colon Cancer	2	190	0.28

Table 1

Mean F1 scores computed across classes and individuals for all 13 considered datasets. No hyperparameter changes were made to the system between each dataset. Please refer to the paper for detailed results and analysis.

Num. planes, $i =$	9	6	3
Single-Planar Ensemble	0.717 ± 0.019	0.714 ± 0.021	0.710 ± 0.024
Multi-Planar U-Net	0.743 ± 0.028	0.737 ± 0.027	0.717 ± 0.030

Table 2

Comparison of the Multi-Planar U-Net and ensembles of 3, 6 and 9 single-planar 2D U-Nets each fit to individual planes on the MICCAI dataset. Shown scores are mean F1.

	MICCAI	HaRP	OAI	MSD T4
3D U-Net w. rotations	0.74 ± 0.04	0.84 ± 0.05	0.81 ± 0.07	0.87 ± 0.04
Multi-Planar U-Net	0.74 ± 0.03	0.85 ± 0.03	0.84 ± 0.07	0.88 ± 0.04

Table 3

Comparison of the Multi-Planar U-Net and 3D U-Net across 4 different datasets. Shown scores are mean F1. The 3D model was trained on isotropic inputs of 64-cube voxels with rotation augmentation. On MSD T4 (Medical Segmentation Decathlon, task 4), the 3D model segments entire images in a single pass.

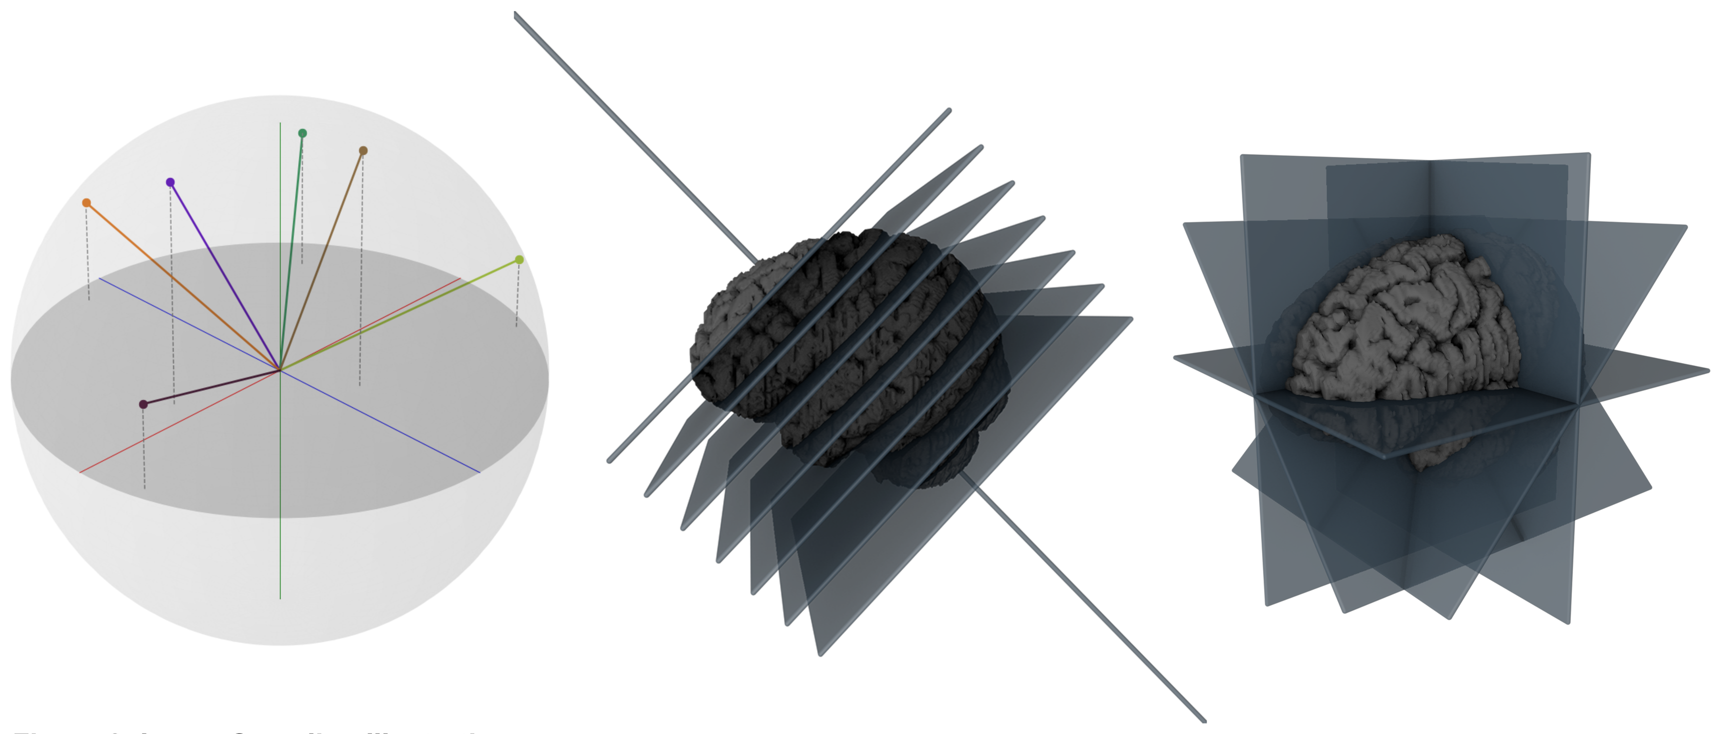

Figure 2: Image Sampling Illustrations

Left: Visualization of a set of sampled planar axis unit vectors. Middle: Visualization of images sampled along one plane. Right: Visualization of multiple images sampled along multiple unique planes.

Open Source Implementation

We provide a regularly maintained, open source implementation of the Multi-Planar U-Net.

https://github.com/perslev/MultiPlanarUNet

References

1. Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. "U-Net: Convolutional Networks for Biomedical Image Segmentation." arXiv [cs.CV]. arXiv. http://arxiv.org/abs/1505.04597.
2. Simpson, Amber L., Michela Antonelli, Spyridon Bakas, Michel Bilello, Keyvan Farahani, Bram van Ginneken, Annette Kopp-Schneider, et al. 2019. "A Large Annotated Medical Image Dataset for the Development and Evaluation of Segmentation Algorithms." arXiv [cs.CV]. arXiv. http://arxiv.org/abs/1902.09063.
3. Isensee, Fabian, Jens Petersen, Andre Klein, David Zimmerer, Paul F. Jaeger, Simon Kohl, Jakob Wasserthal, et al. 2018. "nnU-Net: Self-Adapting Framework for U-Net-Based Medical Image Segmentation." arXiv [cs.CV]. arXiv. http://arxiv.org/abs/1809.10486.
4. Marcus, Daniel S., Tracy H. Wang, Jamie Parker, John G. Csernansky, John C. Morris, and Randy L. Buckner. 2007. "Open Access Series of Imaging Studies (OASIS): Cross-Sectional MRI Data in Young, Middle Aged, Nondemented, and Demented Older Adults." Journal of Cognitive Neuroscience 19 (9): 1498–1507.
5. Boccardi, Marina, Martina Bocchetta, Félix C. Morency, D. Louis Collins, Masami Nishikawa, Rossana Ganzola, Michel J. Grothe, et al. 2015. "Training Labels for Hippocampal Segmentation Based on the EADC-ADNI Harmonized Hippocampal Protocol." Alzheimer's & Dementia: The Journal of the Alzheimer's Association 11 (2): 175–83.
6. Dam, Erik B., Martin Lillholm, Joselene Marques, and Mads Nielsen. 2015. "Automatic Segmentation of High- and Low-Field Knee MRIs Using Knee Image Quantification with Data from the Osteoarthritis Initiative." Journal of Medical Imaging (Bellingham, Wash.) 2 (2): 024001.