Learn2Synth: Learning Optimal Data Synthesis Using Hypergradients for Brain Image Segmentation

Xiaoling Hu

Athinoula A. Martinos Center for Biomedical Imaging

MGH/Harvard Medical School

https://huxiaoling.github.io/

SynthSeg

- **Training networks with entirely synthetic data**

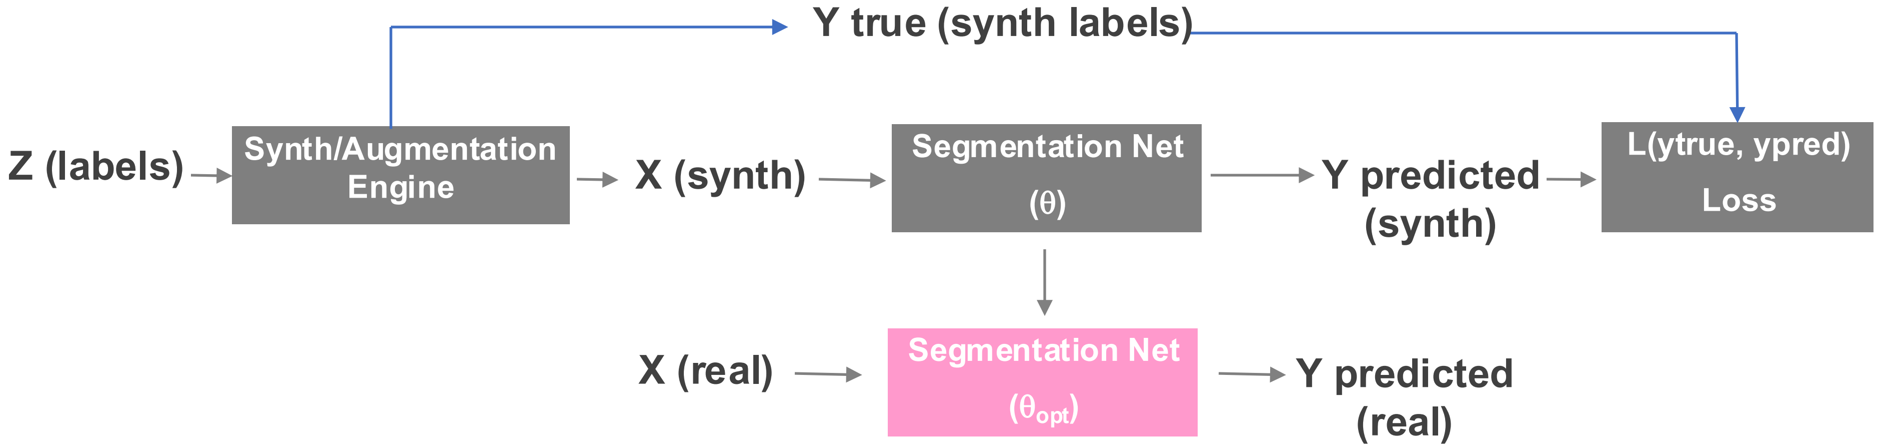

Billot, Benjamin, et al. "SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining." **MedIA** (2023).

Learn2Synth

- **Goal: improve the real world accuracy of a network trained on synthetic data**
  - Bias the synthesis such that the trained network works best on real data
- **Idea**
  - Augment synthetic training examples (generated under ad-hoc rules) by a *trainable* network
    - (Only) Augmented synthetic data used to train seg. network
  - Incite aug. network to generate samples that better train the seg. network: benefit from real labeled examples
  - Don’t care if the synthetic data is realistic or not

Learn2Synth

- **Training strategy**
  - Step1: freeze the aug. network and update the seg. network when running a "synthetic" training step
    - (Only) Augmented synthetic data used to train seg. network
  - Step2: freeze the seg network and update the aug. network when running the "real" training step
    - Incite aug. network to generate samples that better train the seg. network

Pipeline

- **Step1: 'Synthetic' pass and update segmentation branch**
  - Black arrow: data flow; red arrow: gradient flow

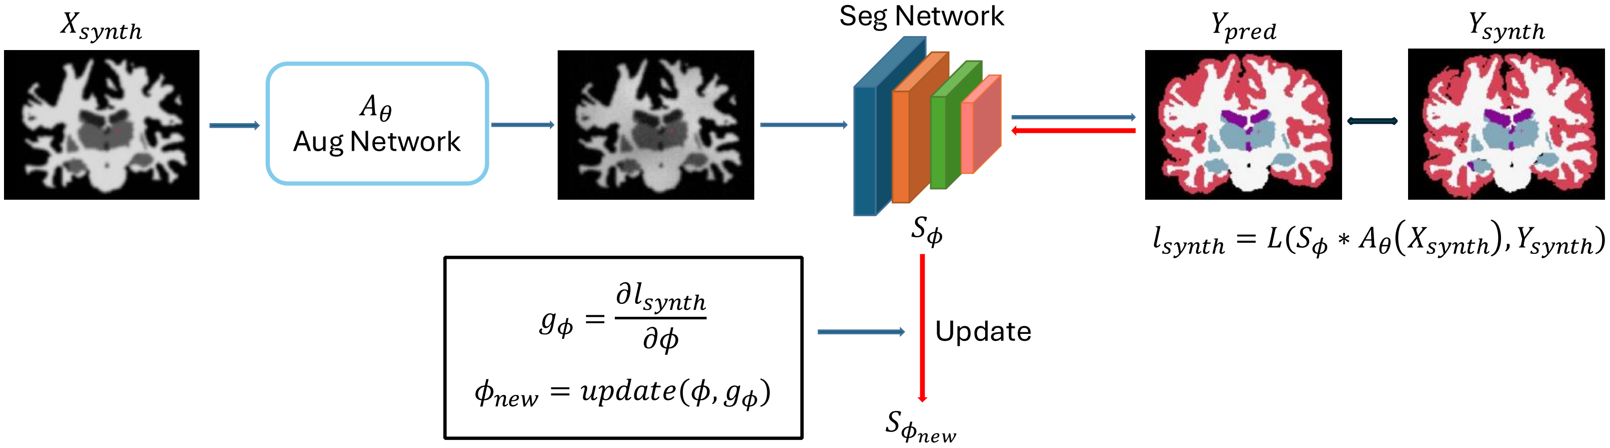

Pipeline

- **Step2: 'Real' pass and update augmentation network**
  - Black arrow: data flow; red arrow: gradient flow

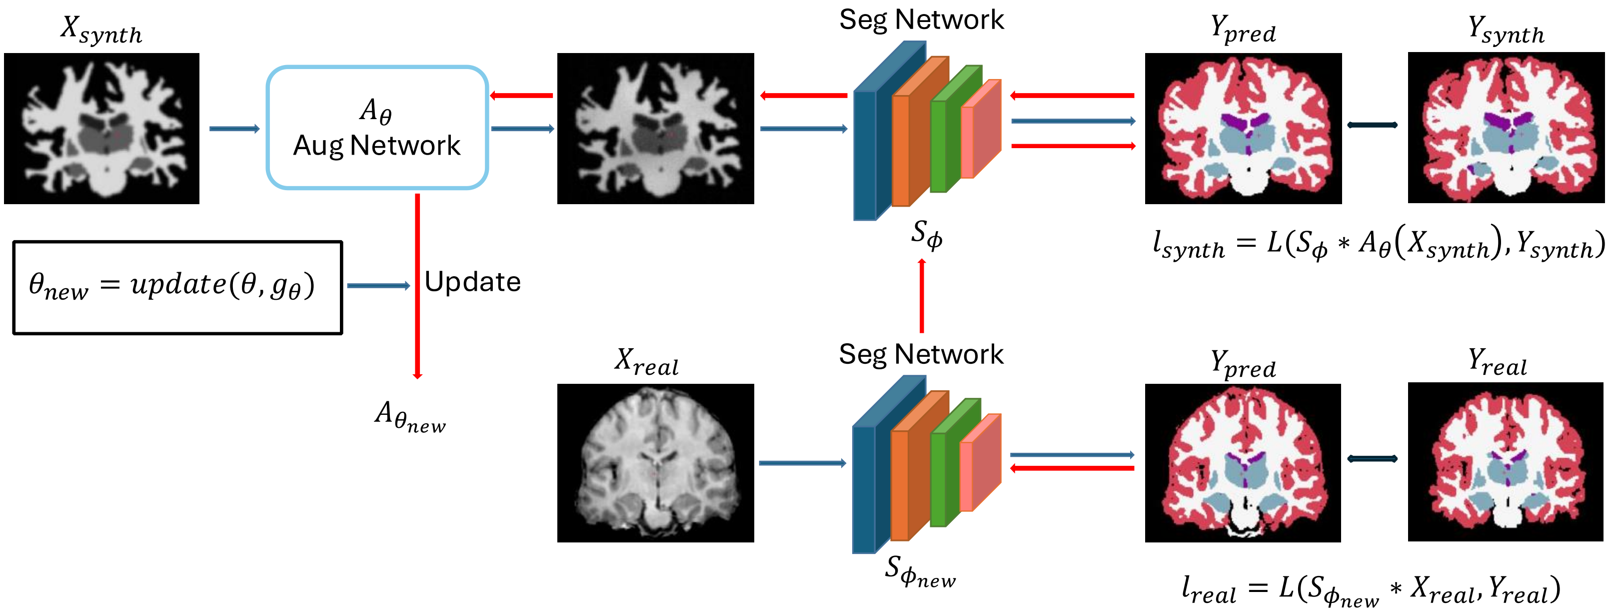

Framework

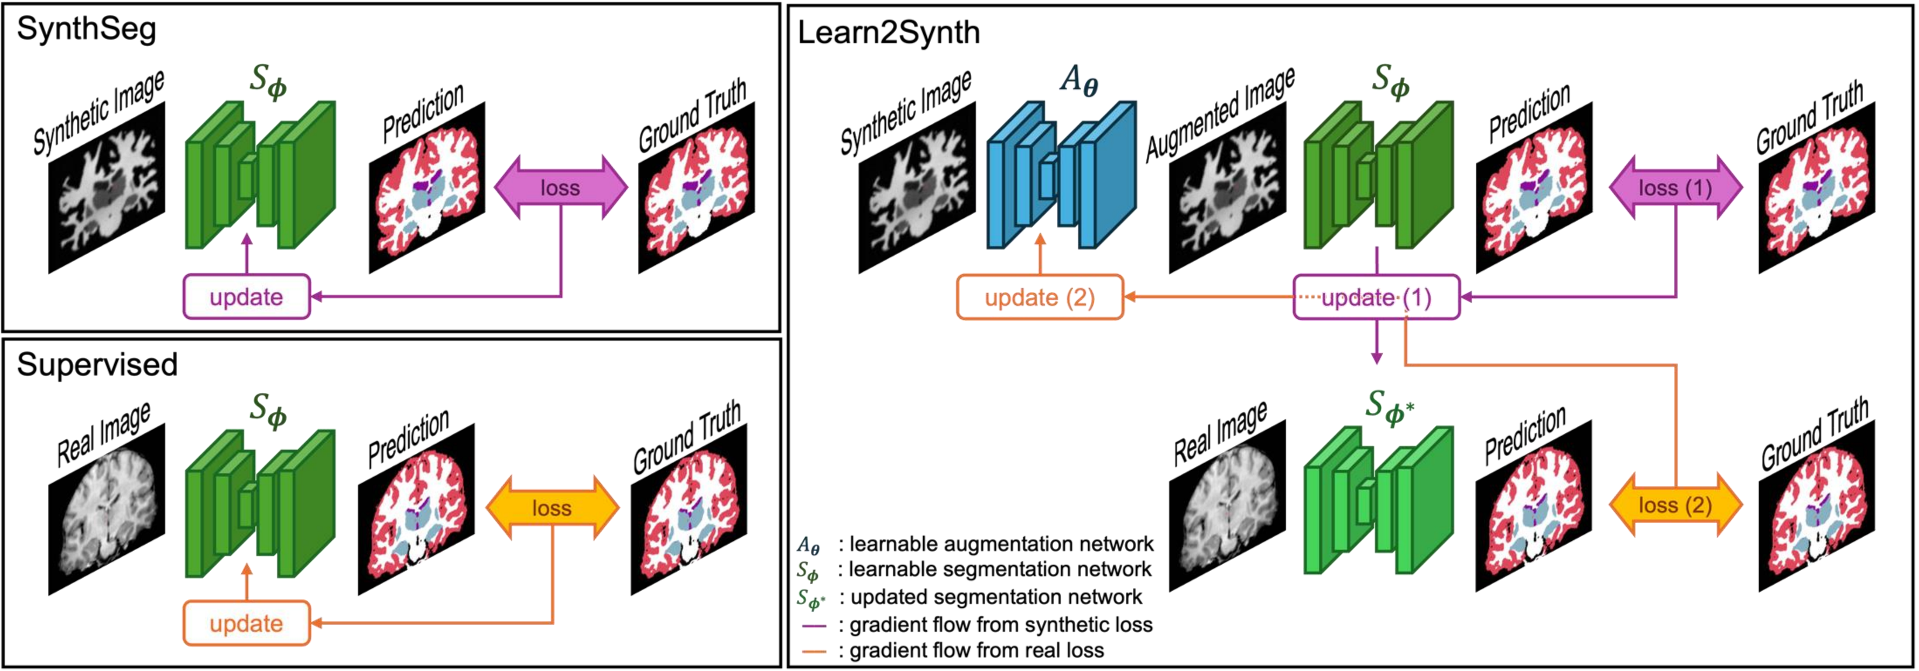

SynthSeg
Synthetic Image
S_ϕ
Prediction
Ground Truth
loss
update
Supervised
Real Image
S_ϕ
Prediction
Ground Truth
loss
update
Learn2Synth
Synthetic Image
A_θ
Augmented Image
S_ϕ
Prediction
Ground Truth
loss (1)
update (2)
update (1)
Real Image
S_{ϕ^*}
Prediction
Ground Truth
loss (2)
A_θ : learnable augmentation network
S_ϕ : learnable segmentation network
S_{ϕ^*} : updated segmentation network
: gradient flow from synthetic loss
: gradient flow from real loss

Results on synthetic datasets

- **Settings**
  - Train both naïve SynthSeg and Learn2Synth with different noise levels added to the training images
    - $\boldsymbol{\sigma}$ = 0, 0.05, 0.1, 0.15 and uniformly sampled from [0.05, 0.2]
    - Test on images with corresponding noise levels

Quantitative results

- Check if the network is learning what we expect

Preset $\hat{\sigma}$	0	0.050	0.100	0.150	[0.025 0.2]
Inferred σ^*	0.001	0.042	0.098	0.146	0.134

Table 1. Inferred noise σ under the “noise-only” setting.

Qualitative results

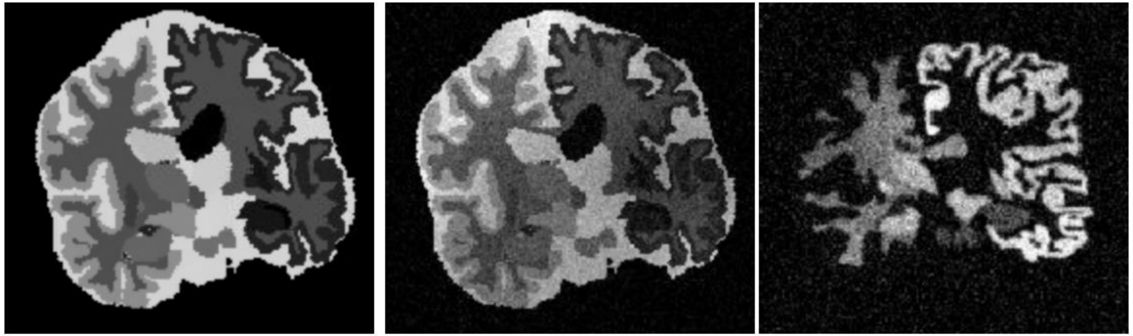

(a) Noise-free image (b) Learned image (c) Real image

Figure 3. Illustration of learned synthesized image. (a) Noise-free image, (b) learned synthesized image, and (c) real image.

Quantitative results

- Segmentation performance

Train \ Test	Model	$\sigma = 0$	$\sigma = 0.05$	$\sigma = 0.1$	$\sigma = 0.15$	$\sigma \sim U$
$\sigma = 0$	N	0.913	0.310	0.209	0.142	0.205
	L	**0.923**	0.346	0.172	0.092	0.151
	O	0.919	0.361	0.189	0.105	0.167
$\sigma = 0.05$	N	0.889	0.874	0.718	0.520	0.646
	L	0.900	0.887	0.744	0.492	0.635
	O	0.898	**0.890**	0.736	0.508	0.651
$\sigma = 0.1$	N	0.867	0.858	0.841	0.772	0.815
	L	0.874	0.873	**0.858**	0.799	0.831
	O	0.871	0.869	0.847	0.810	0.829
$\sigma = 0.15$	N	0.864	0.857	0.847	0.829	0.838
	L	0.860	0.855	0.854	**0.840**	0.834
	O	0.867	0.860	0.852	0.836	0.841
$\sigma \sim U$	N	0.870	0.867	0.850	0.826	0.831
	L	0.866	0.864	0.851	0.833	0.839
	O	0.872	0.870	0.853	0.837	**0.843**

Table 2. Segmentation accuracy (Dice) for naive SynthSeg (N), *Learn2Synth* (*L*) and Optimized SynthSeg (O).

Results on real-world datasets

- 2D Segmentation performance

Method	ABIDE	OASIS3
Supervised UNet	0.908	0.899
SAMSEG	0.875	0.841
Naive SynthSeg	0.869	0.831
Mixed SynthSeg	0.875	0.854
Finetuned SynthSeg	0.871	0.847
AdvChain	0.867	0.848
Learn2Synth (parametric setting fixed σ)	0.878	0.860
Learn2Synth (parametric setting varying σ)	**0.881**	0.857
Learn2Synth (nonparametric setting fixed σ)	0.879	0.875
Learn2Synth (nonparametric setting varying σ)	0.874	**0.881**

Table 5. Segmentation accuracy for *Learn2Synth* and the baselines on ABIDE and OASIS3.

Results on real-world datasets

- Generalizability

Setting	# of Train, Val., Test	Test Set (MPRAGE)	3°	5°	20°	30°
SynthSeg	/	0.861 ± 0.028	0.776 ± 0.012	0.694 ± 0.027	0.766 ± 0.018	0.781 ± 0.016
Supervised UNet	29, 5, 5	$\mathbf{0.941} \pm 0.002$	0.419 ± 0.025	0.396 ± 0.020	0.671 ± 0.026	0.769 ± 0.022
Supervised UNet	5, 5, 29	0.907 ± 0.007	0.397 ± 0.018	0.413 ± 0.016	0.586 ± 0.033	0.692 ± 0.033
Supervised UNet	1, 5, 33	0.885 ± 0.010	0.267 ± 0.019	0.247 ± 0.013	0.544 ± 0.026	0.651 ± 0.025
Learn2Synth	29, 5, 5	0.895 ± 0.011	$\mathbf{0.804} \pm 0.014$	$\mathbf{0.789} \pm 0.010$	0.785 ± 0.025	0.797 ± 0.023
Learn2Synth	5, 5, 29	0.867 ± 0.030	0.798 ± 0.013	$\mathbf{0.789} \pm 0.010$	$\mathbf{0.795} \pm 0.026$	$\mathbf{0.799} \pm 0.024$
Learn2Synth	1, 5, 33	0.718 ± 0.045	0.606 ± 0.020	0.603 ± 0.017	0.690 ± 0.036	0.707 ± 0.032

Table 6. Performance comparison of different models across various datasets.

Results on real-world datasets

- Generalizability

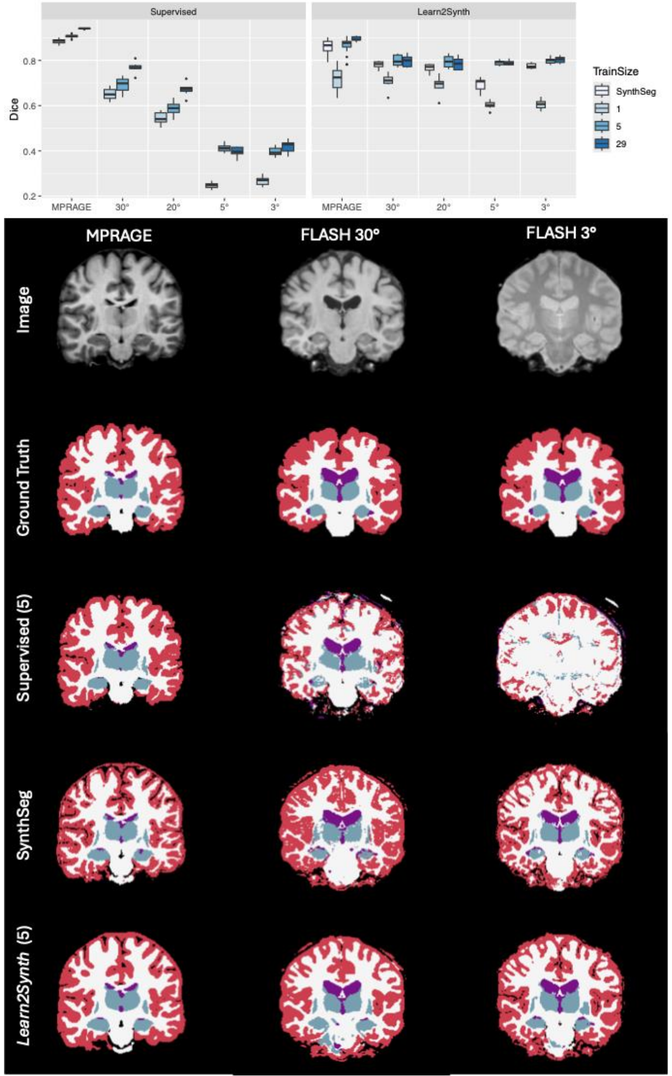

Thank you for your attention!

Q&A