NiftyNet

An open-source convolutional neural networks platform for research in medical image analysis and image-guided therapy

MICCAI 2018
Tutorial on Tools Allowing Clinical Translation of Image Computing ALgorithms [T.A.C.T.I.C.AL.]

Presented by Tom Vercauteren

http://niftynet.io
https://github.com/NifTK/NiftyNet

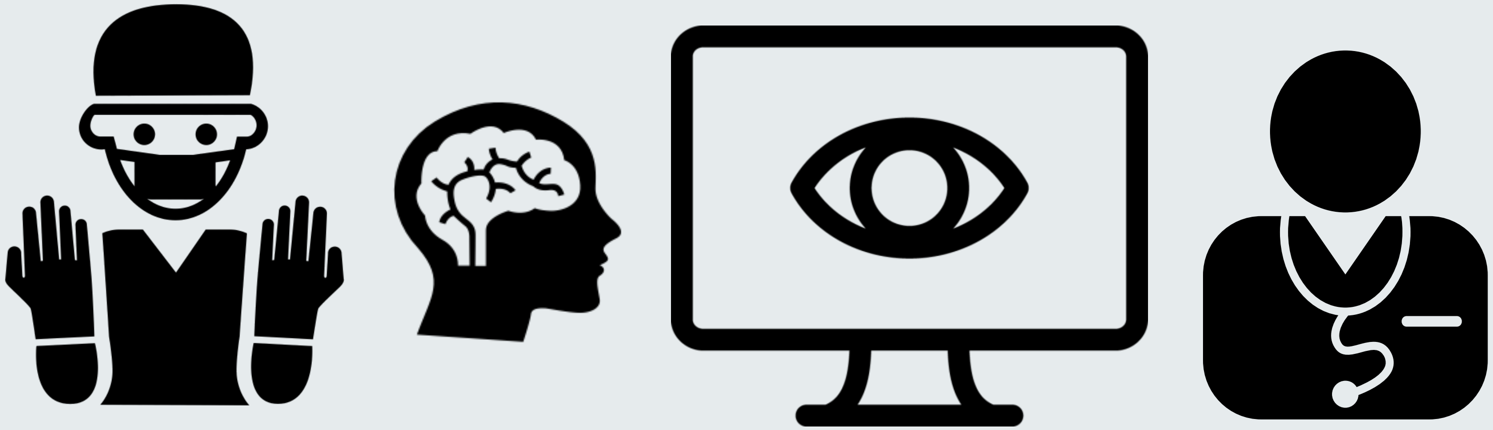

EPSRC

wellcome

KING'S College LONDON

NiftyNet: An open consortium for deep learning in medical imaging


NiftyNet: a deep-learning platform for medical imaging

Abstract

Infrastructure paper

Gibson, E., Li, W., Sudre, C., Fidon, L., Shakir, D. I., Wang, G., … Vercauteren, T. (2018). NiftyNet: a deep-learning platform for medical imaging. *Computer Methods and Programs in Biomedicine*, *158*, 113–122. doi:10.1016/j.cmpb.2018.01.025. arXiv:1709.03485

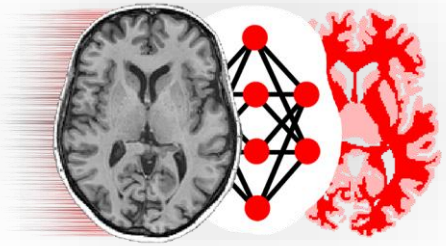

- Apache-2.0 licensed
- Collection of domain specific knowledge and best practices
- Validated, tested, SOTA network/loss/sampler/aggregator implementations
- Best-practice API usage for data management, multiple-GPU and tensorboard support
- Platform for model evaluation, dissemination and deployment

The need for NiftyNet

Relevance of deep learning and idiosyncrasies

Radio/patho/dermato/ophtalmo-logists replaced by AI?

"If you work as a radiologist you're like the coyote that's already over the edge of the cliff. ***People should stop training radiologists now****, it's just completely obvious that in five years deep learning is going to do better than radiologists.*
It might be ten years"

Geoffrey Hinton in 2016

More nuanced opinions

"In the not so far future, machine learning will play a central role in radiology, becoming part of routine workflow and providing daily real-time clinical diagnostic support.
I predict that within 10 years ***no medical imaging study will be reviewed by a radiologist until it has been pre-analyzed*** *by a machine."*

Nick Bryan in 2016
(past RSNA president)

Representation power of neural networks (NN)

- NN with one hidden layer can represent any bounded continuous function
  - Universal approximation theorem (Cybenko 1989)
- Deep neural networks can represent complex functions more efficiently
- Convolutional neural networks allows exploiting image structure

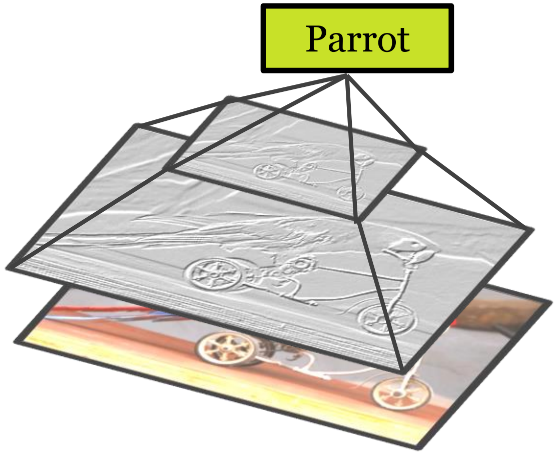

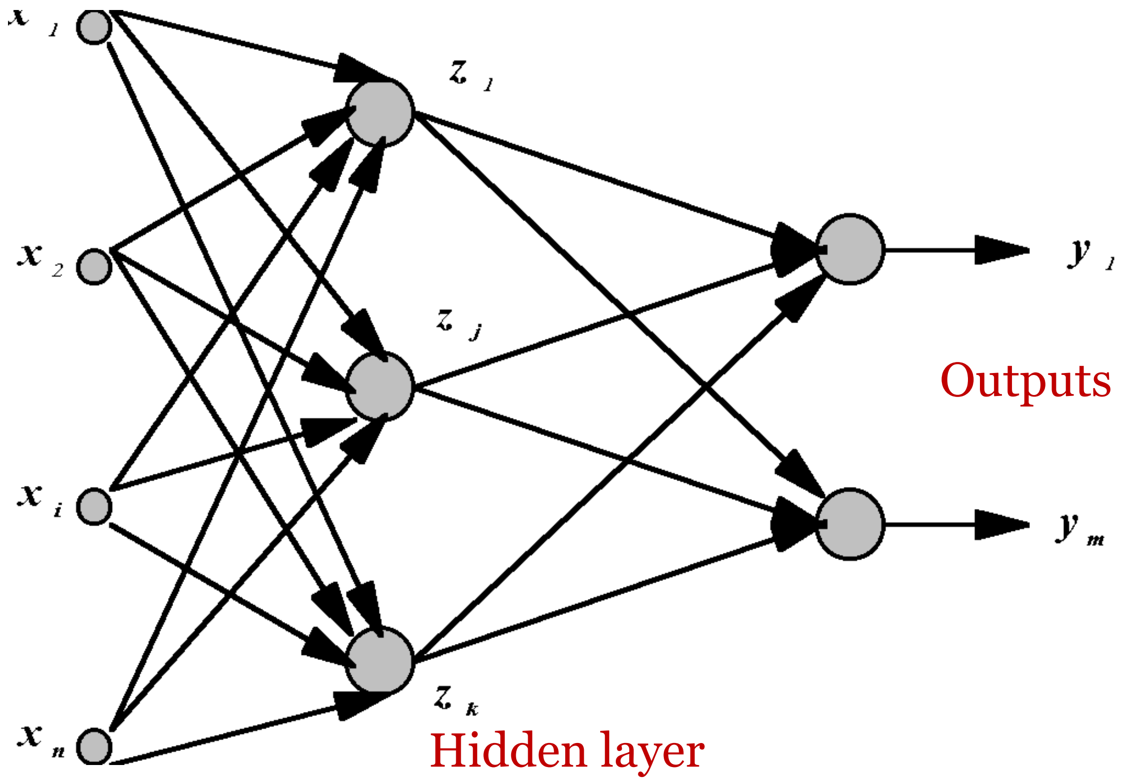

Inputs

$\boldsymbol{x} = (x_1, \dots, x_{784})^T$

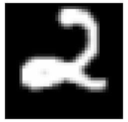

Hidden layer

$z_j = g(\boldsymbol{w}_j^T \boldsymbol{x})$

$g(t) = \dfrac{1}{1 + e^{-t}}$

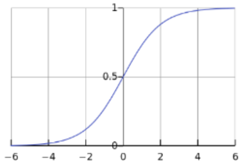

Activation

Slide adapted from Linda Shapiro CSE 455

Tangible successes in medical imaging beyond the hype?

- **As of 2016**, Samsung is using deep learning for breast lesion analysis in its ultrasound product line (S-Detect for Breast)
- Samsung reported using a database of over 10,000 images
- Fair bit of interactivity still involved

Lots of research opportunities to demonstrate impact in medical imaging

- **Sep 2018**
- 3D U-Net for tissue segmentation
  - Training: 877 OCTs
- 3D DenseNet for classification
  - Input: segmentation results
  - Training: 14,884 3D OCTs
- Testing: 997 patients
- doi:10.1038/s41591-018-0107-6

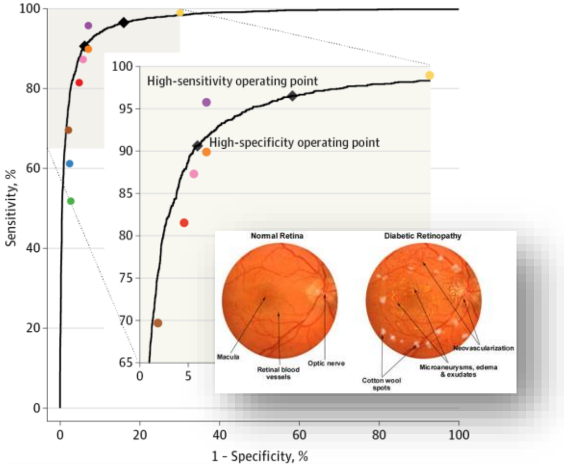

- **Dec 2016**
- 140,000 retinal images
- Ophtalmologist-level performance to detect diabetic retinopathy
- Inception v3, pre-trained on ImageNet
- doi:10.1001/jama.2016.17216

The role of ConvNet toolkits in medical imaging

- Many challenges for deep learning in medical imaging
  - 2D+t, 3D, 4D, …
  - Making use of metadata and prior knowledge: voxel sizes, quantitative intensities, anatomical coordinates, orientation-dependence, etc.
  - Interventional applications
  - Efficient use of computational resources
  - Strictly regulated environment for clinical translation
- Resources should be focused on advancing the field while taking advantage of rapidly evolving techniques
- Reproducibility should be a given
- Translation should benefit from validated approaches

Yet another toolkit?

- Is an I/O library + deep learning not enough?

SimpleITK

NiBabel

+

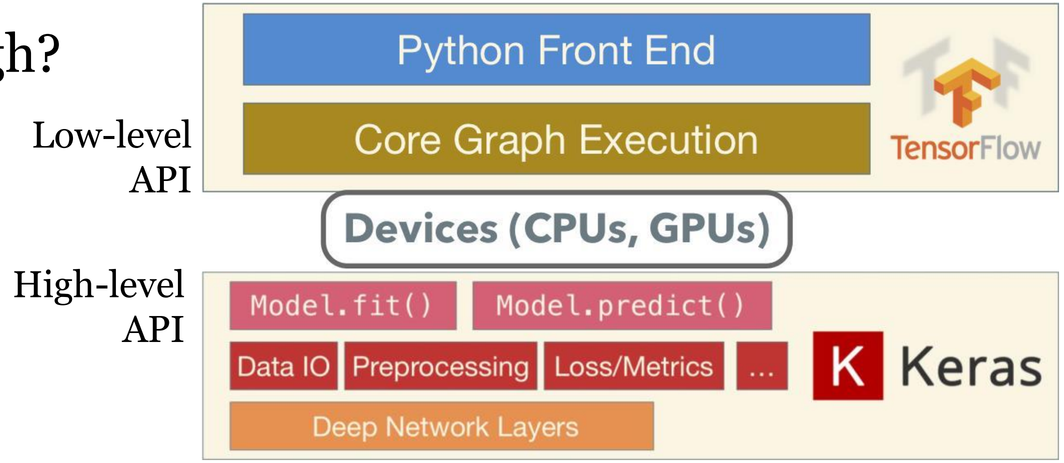

Microsoft CNTK

mxnet

PYTORCH

...

- Need for domain knowledge
  - How to handle 3D, 4D, 5D data and fit it in memory?
  - How to encode voxel-size/scale and coordinate frames intrinsically in deep learning?
  - How to augment medical data in a physics/biology correct manner?
  - Which architectures/loss functions are relevant in medical domain?
  - How to create general best-practice tools that can scale to multiple problems?
  - How do we standardise, compare and share models, and improve their deployability?

Deep learning project routines

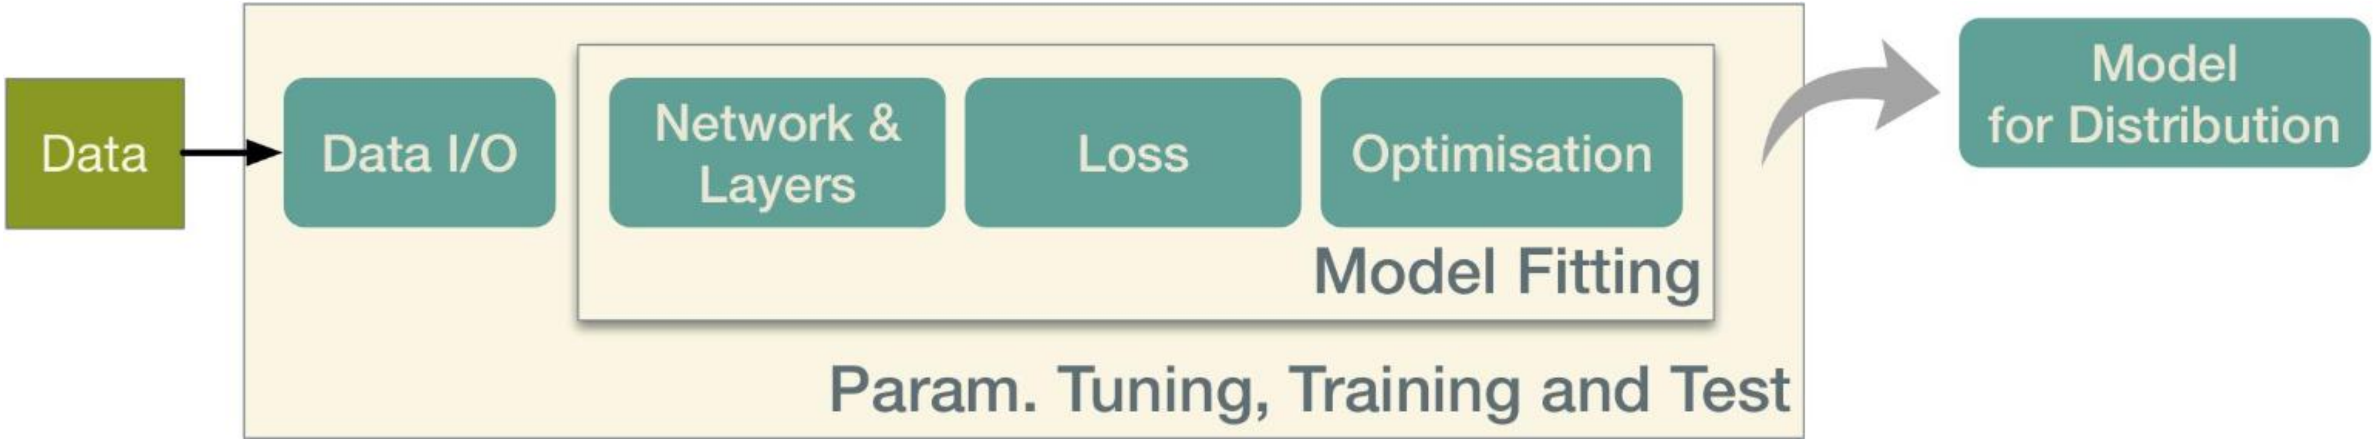

Deep learning in medical imaging – The need for sampling

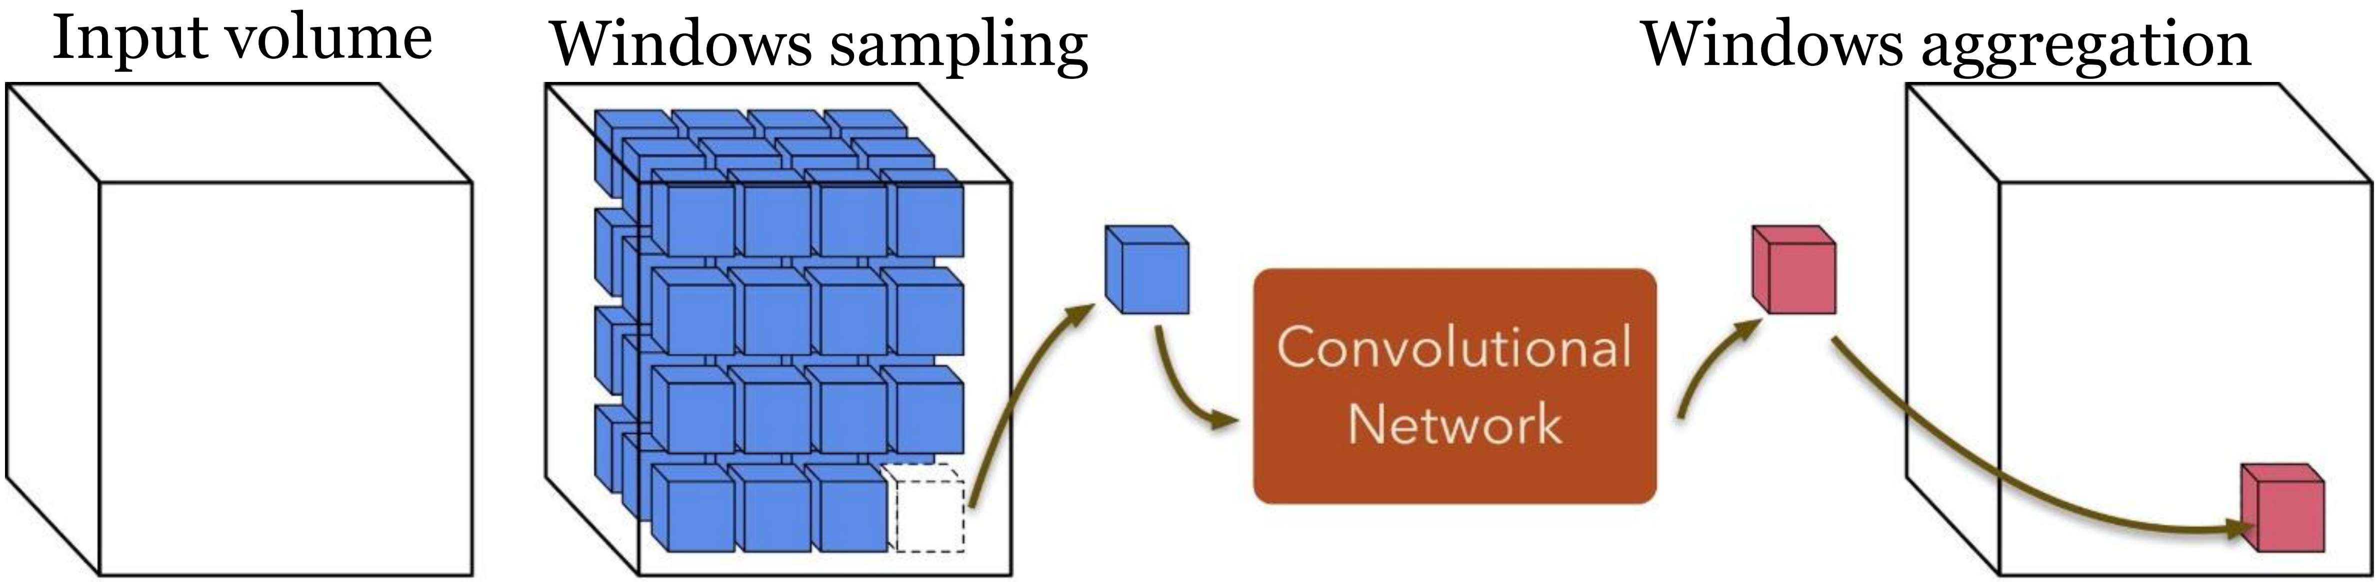

Deep learning in medical imaging – The need for sampling

Input volume

Windows sampling

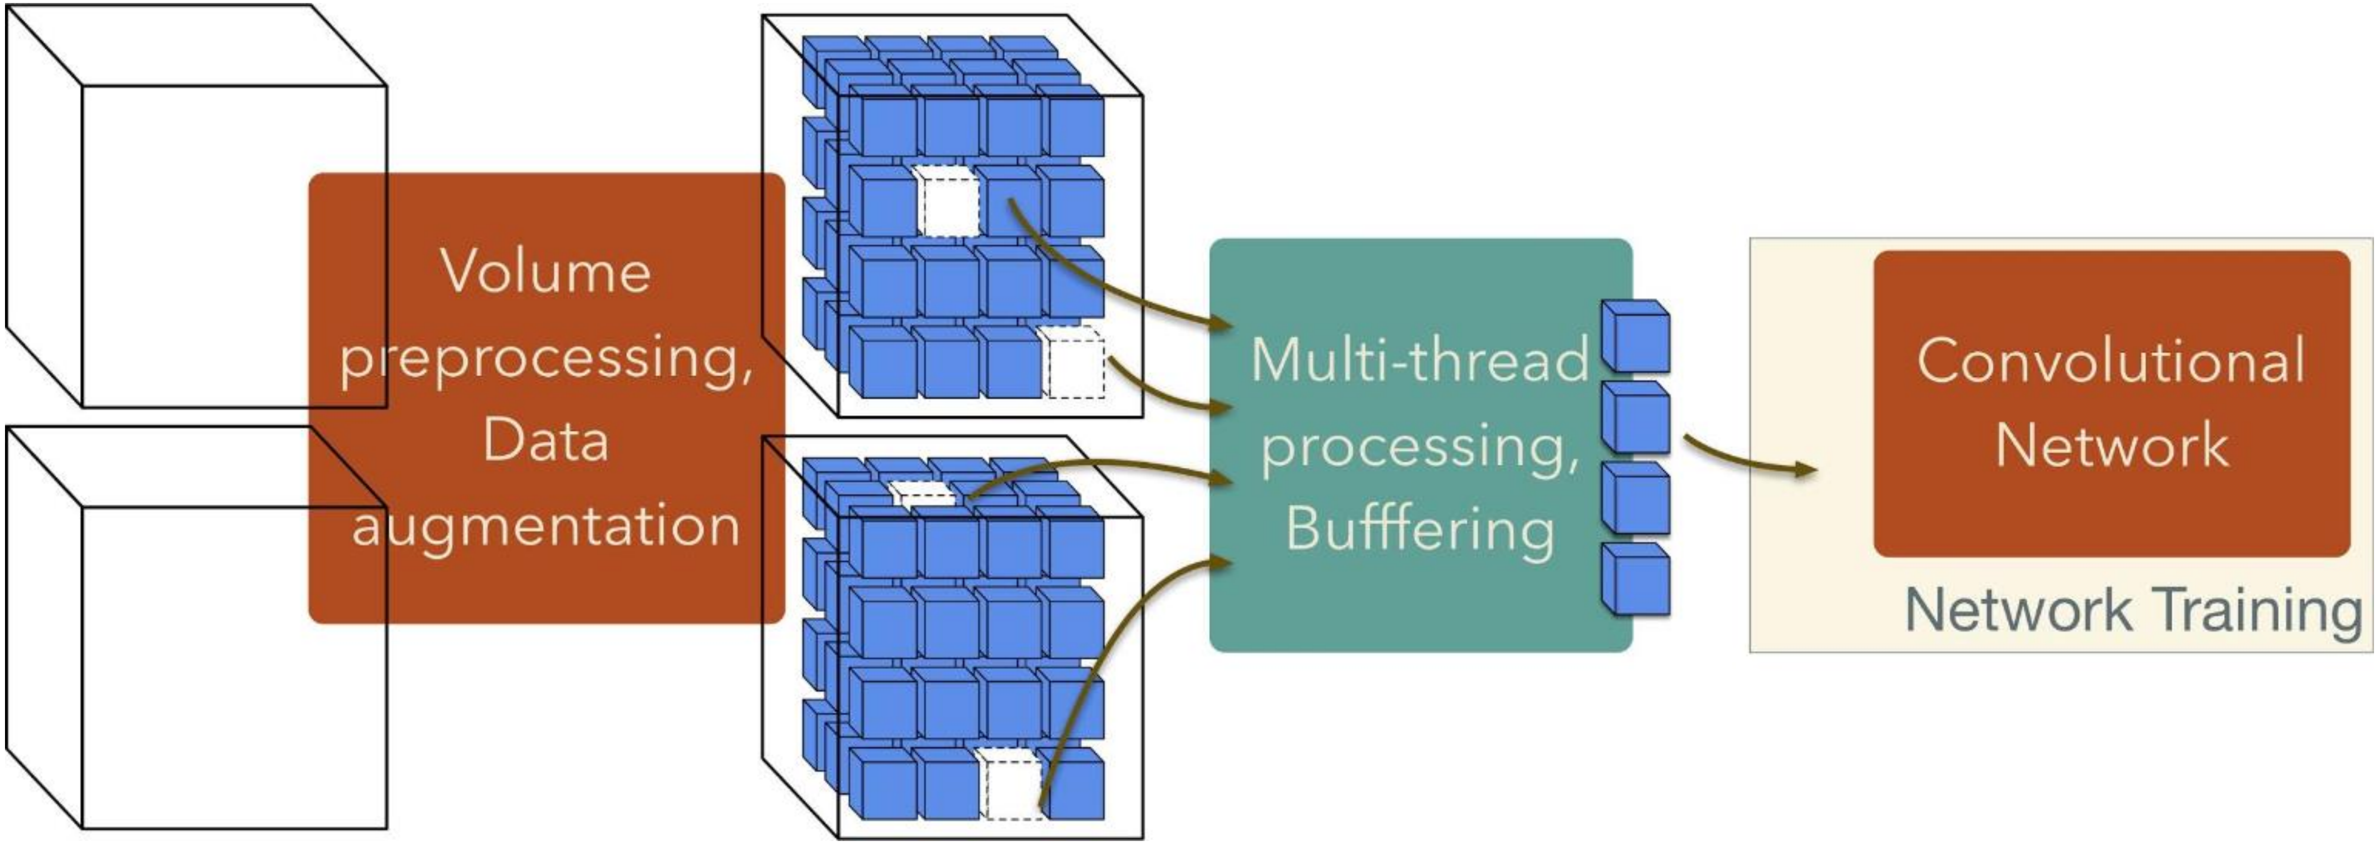

NiftyNet helps fill the gaps

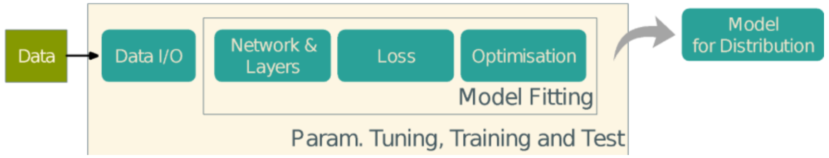

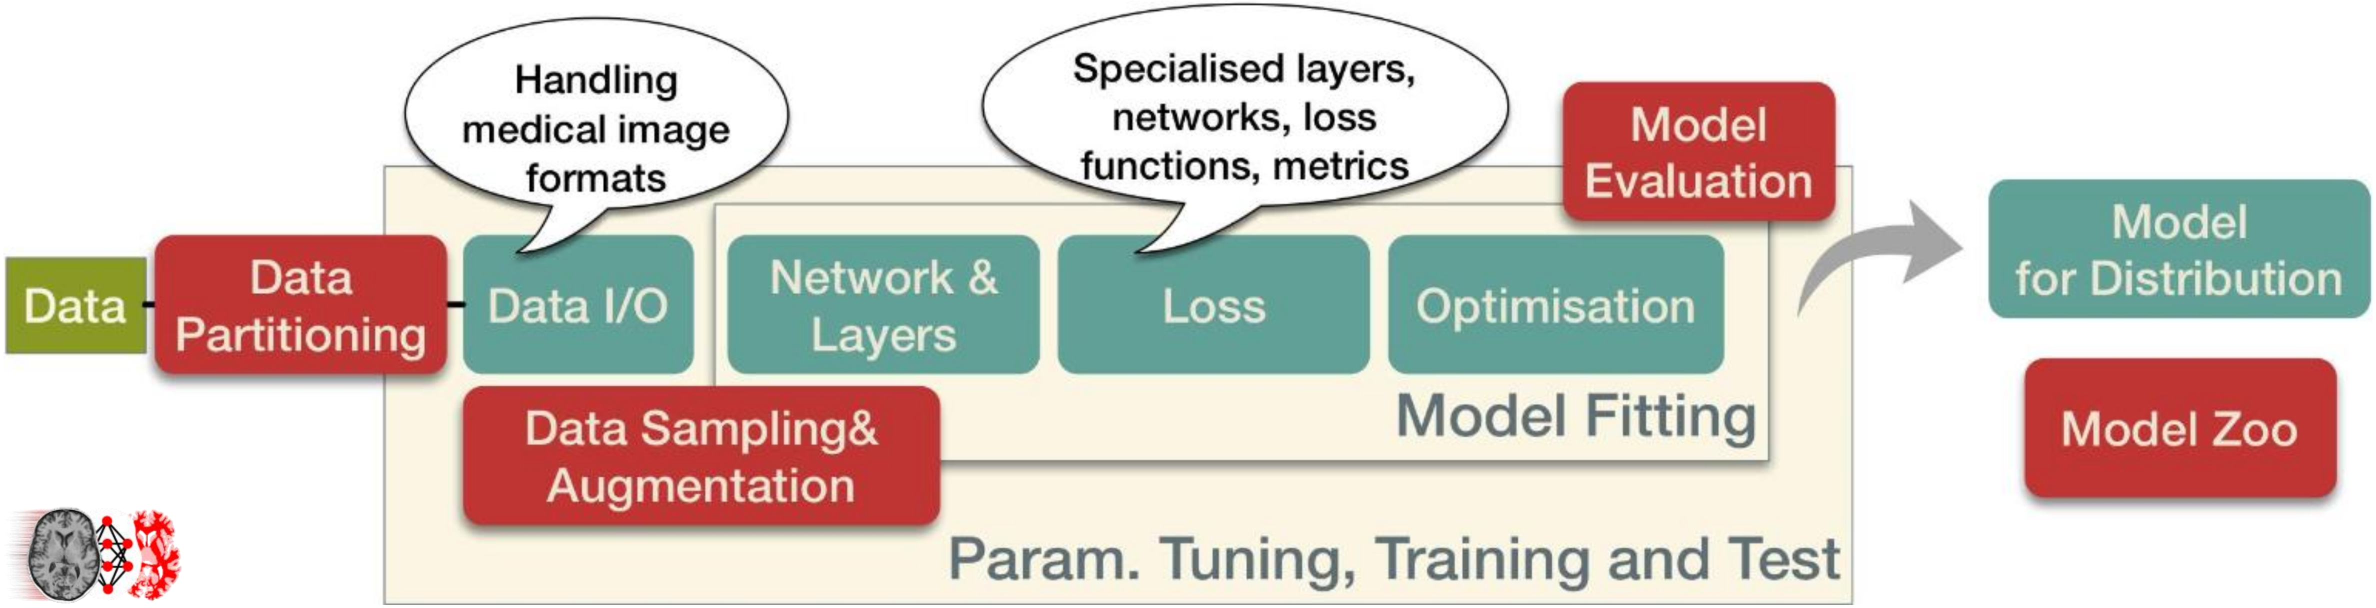

MR to CT synthesis example

Foregoing the need for CT in PET-MR imaging

MR to CT synthesis – An image regression example

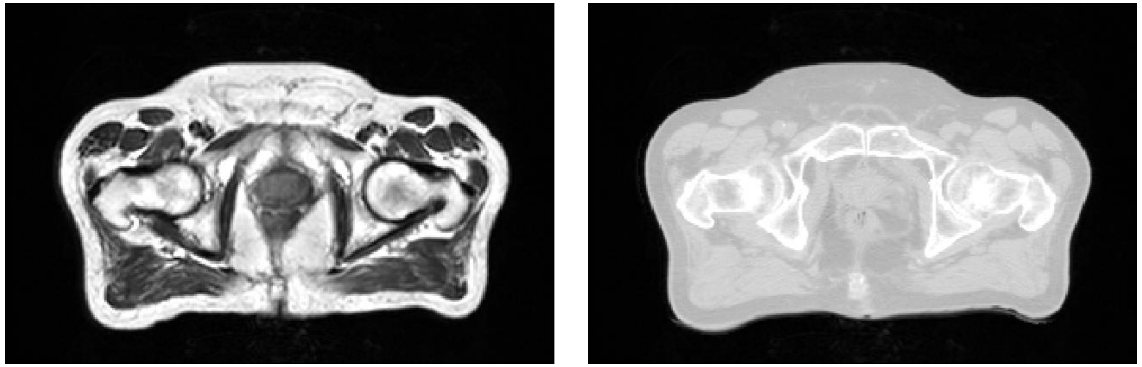

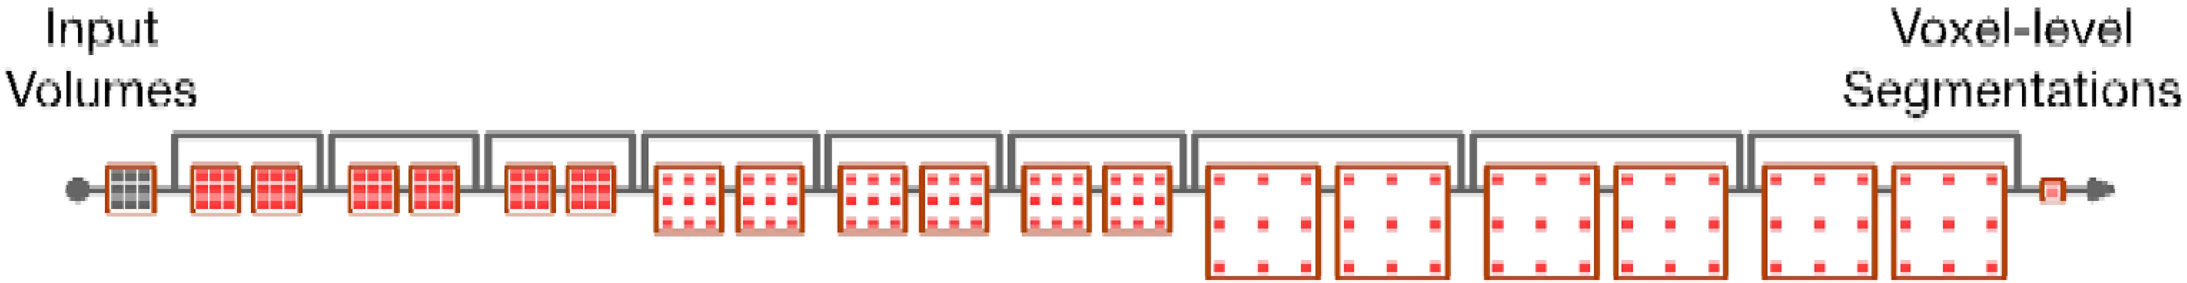

MR to CT synthesis – Sampling at inference

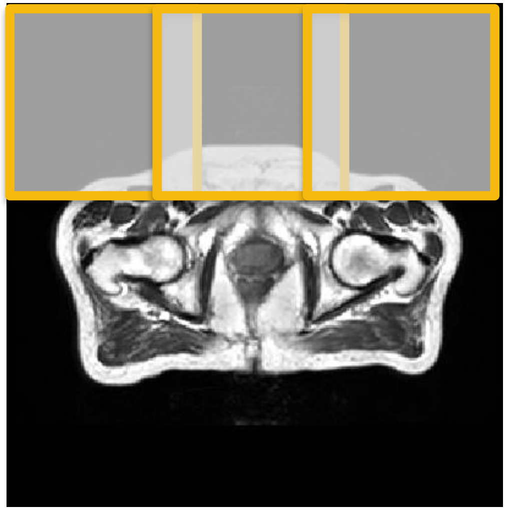

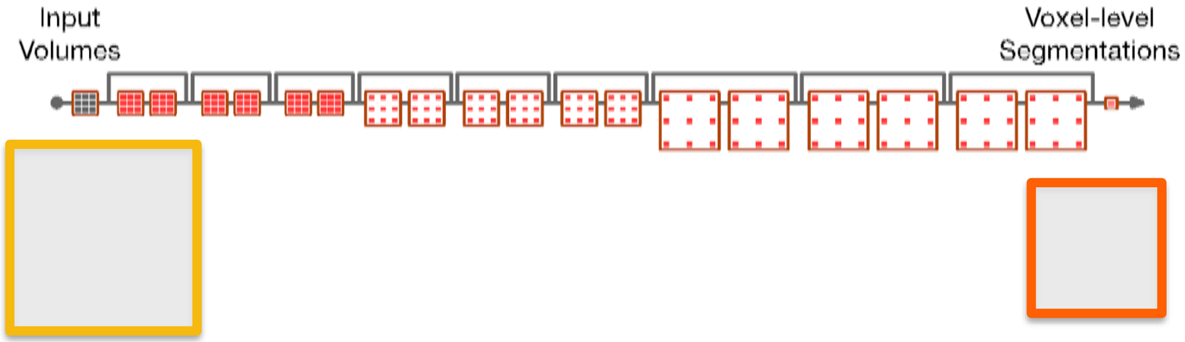

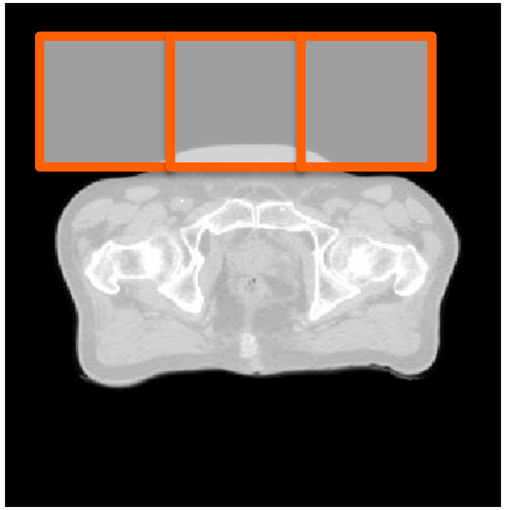

- data augmentation
- window sampling/aggregation
- loss function
- optimisation methods

MR to CT synthesis – Application overview

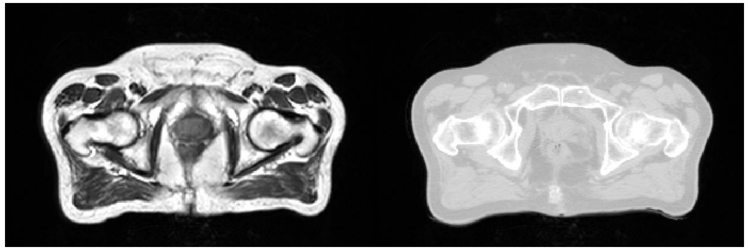

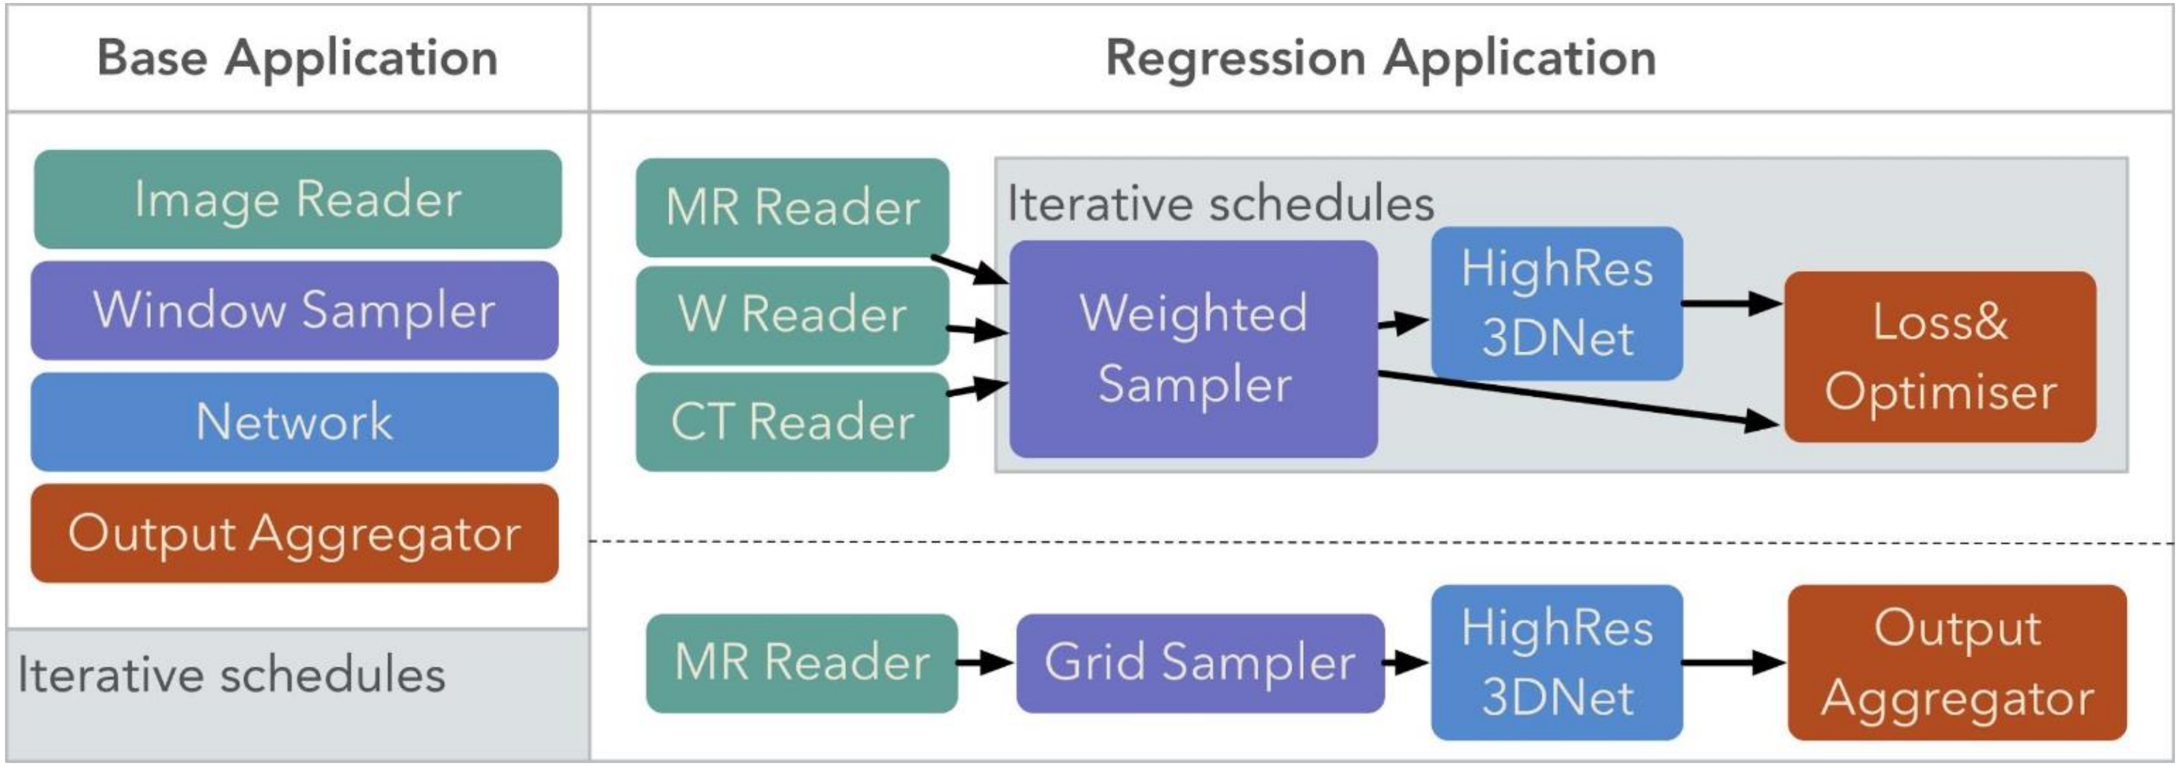

MR to CT synthesis – Added implementation complexity

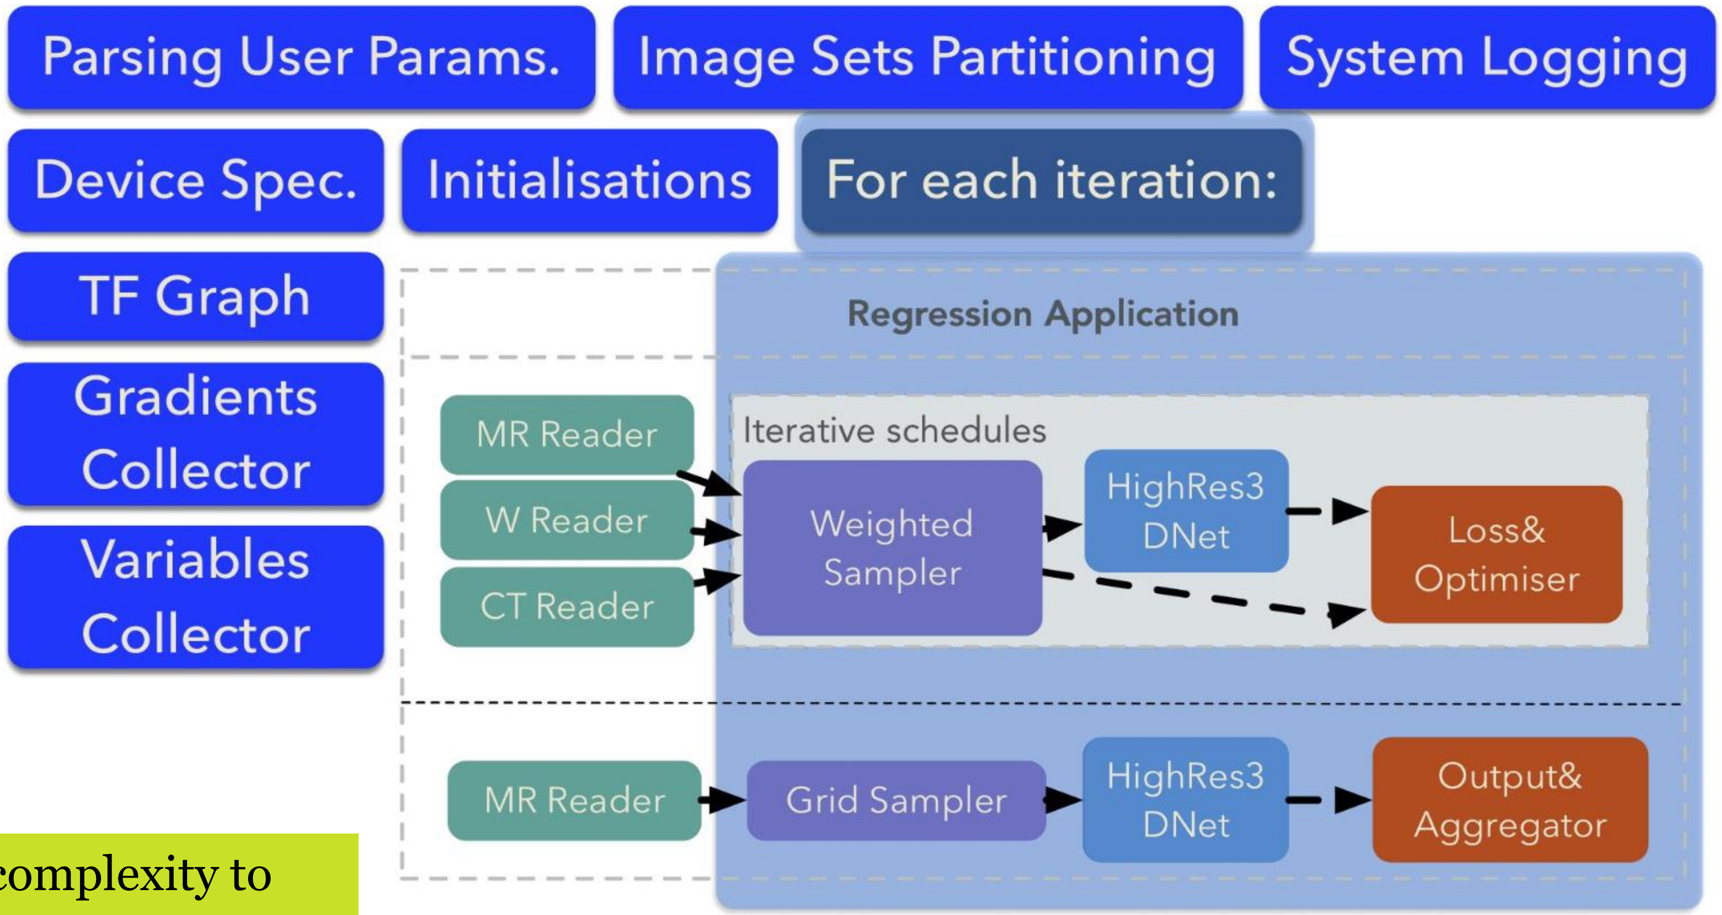

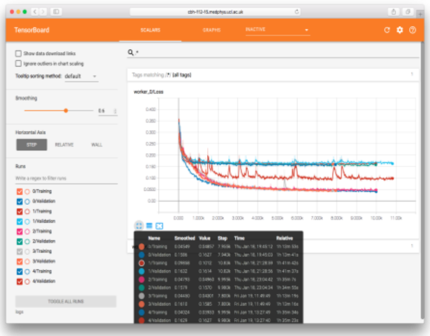

Hiding the complexity to the users

- Config files hierarchy
- Command-line overrides

```
> net_regress train -c my_config.ini --decay 1e-5 --lr 1e-3
```

MR to CT synthesis – Advanced algorithms

Berger et al., arXiv: 1709.02764

Weighted sampling
Error maps from the previous iterations as sampling weights

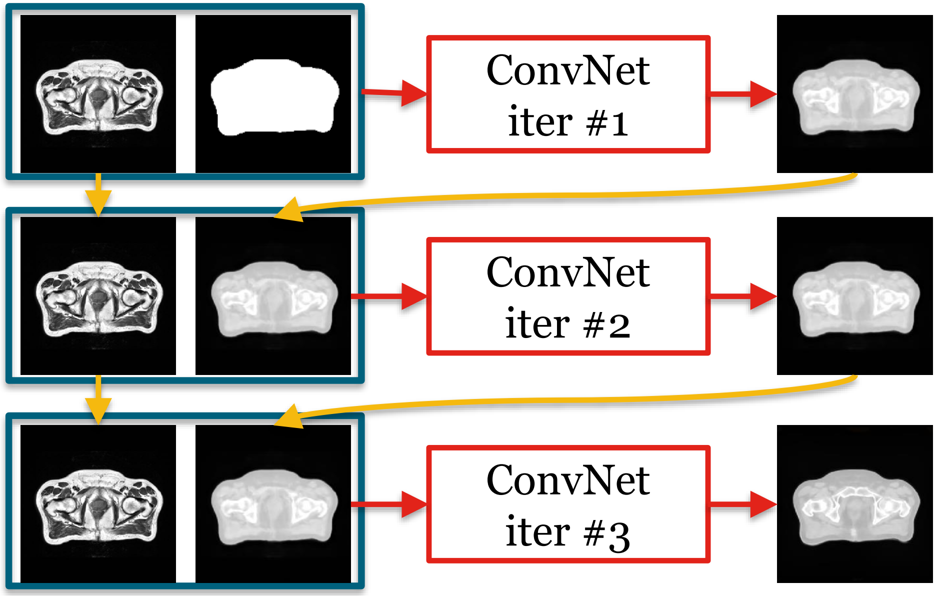

"Autocontext"

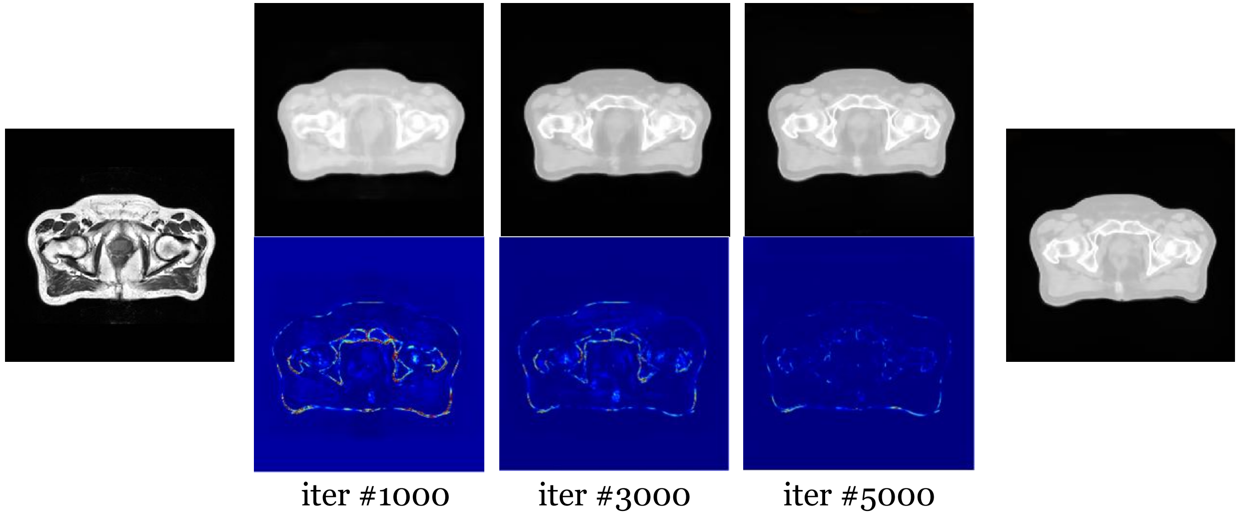

Beyond regression/segmentation/classification – MR to US registration

Hu et al., MedIA 21018 and arXiv: 1711.01666 (2018).

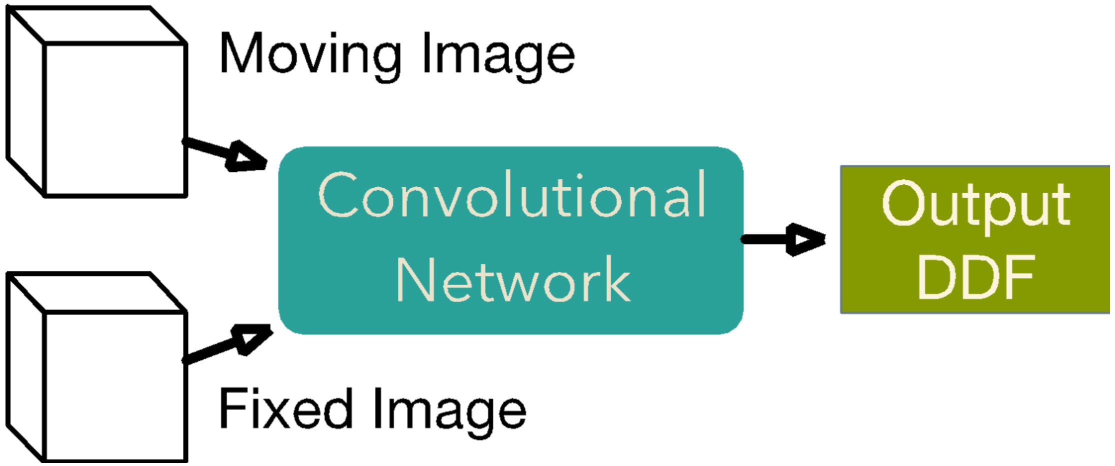

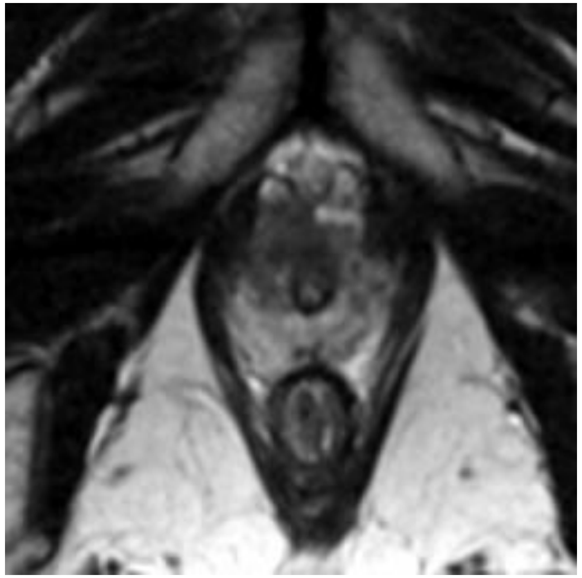

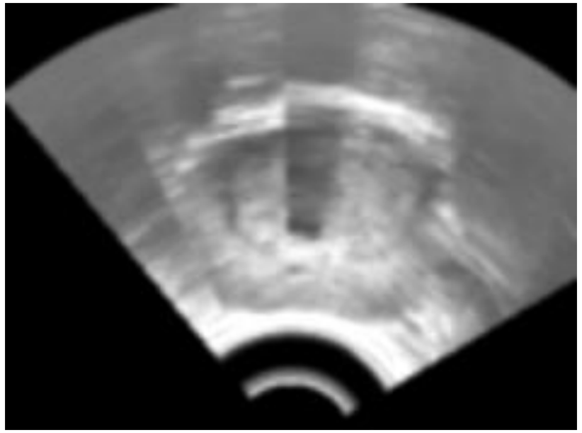

Beyond regression/segmentation/classification – MR to US registration

Hu et al., MedIA 21018 and arXiv: 1711.01666 (2018).

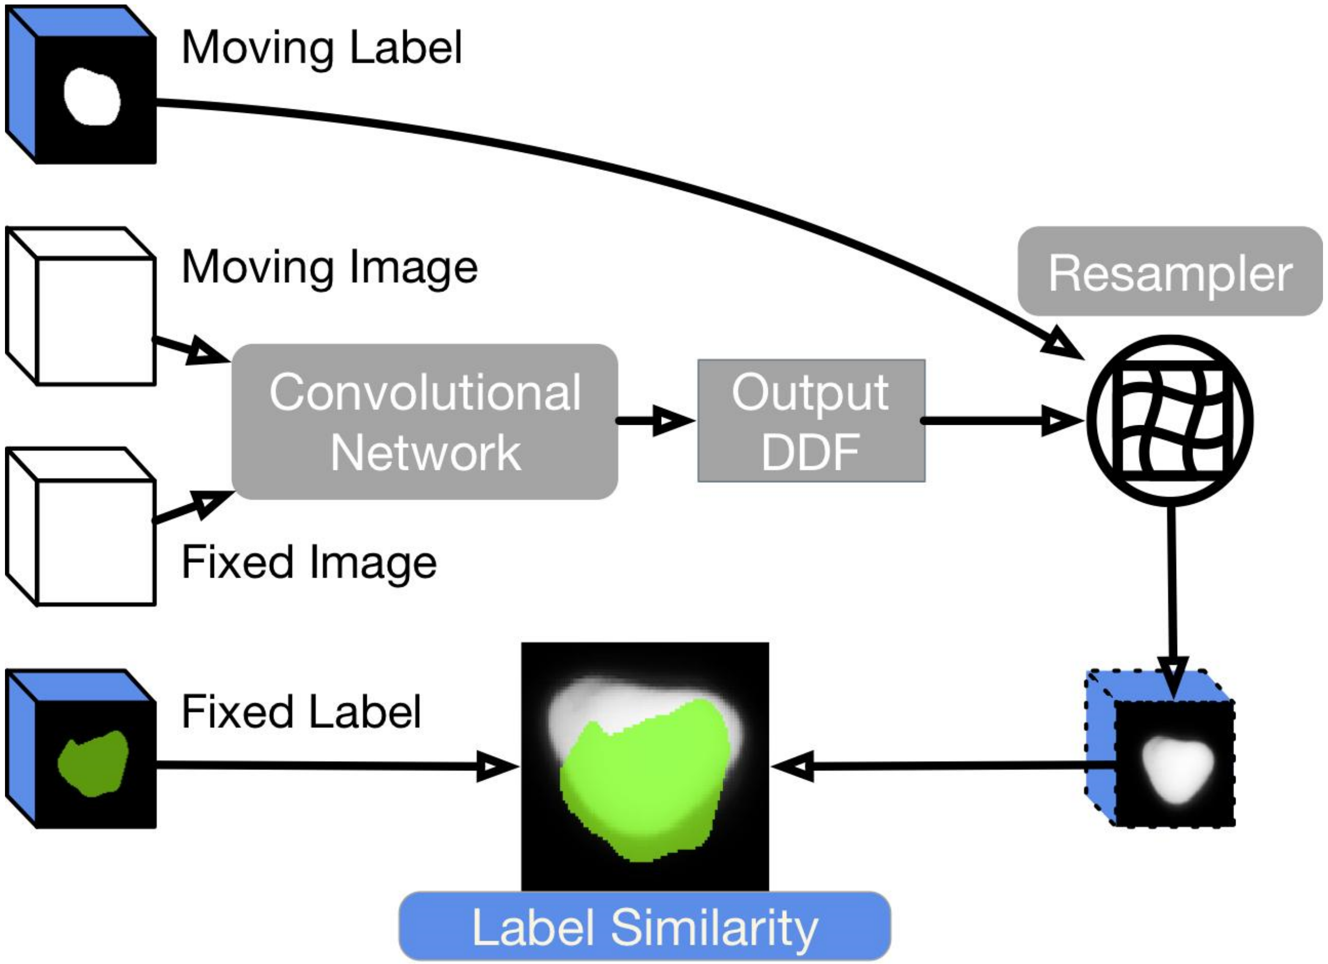

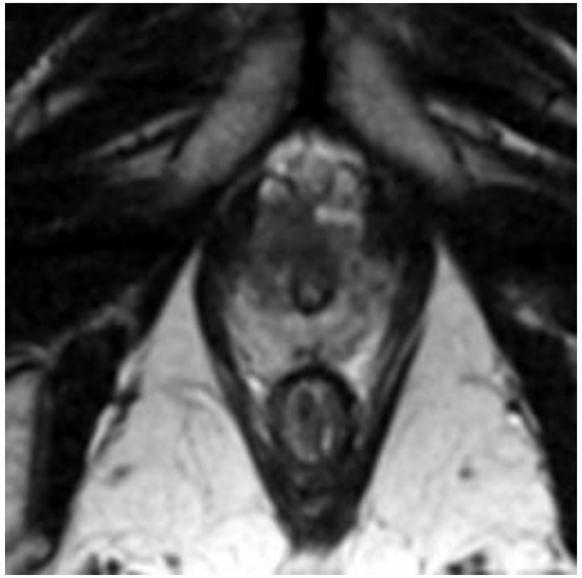

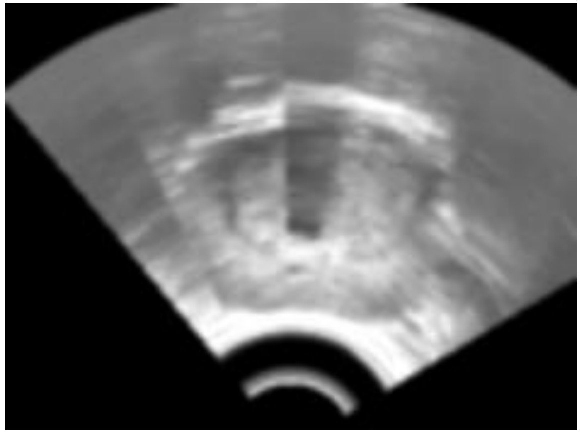

Designing new applications with NiftyNet

```
> net_run train -c my_config.ini -a registration_app
```

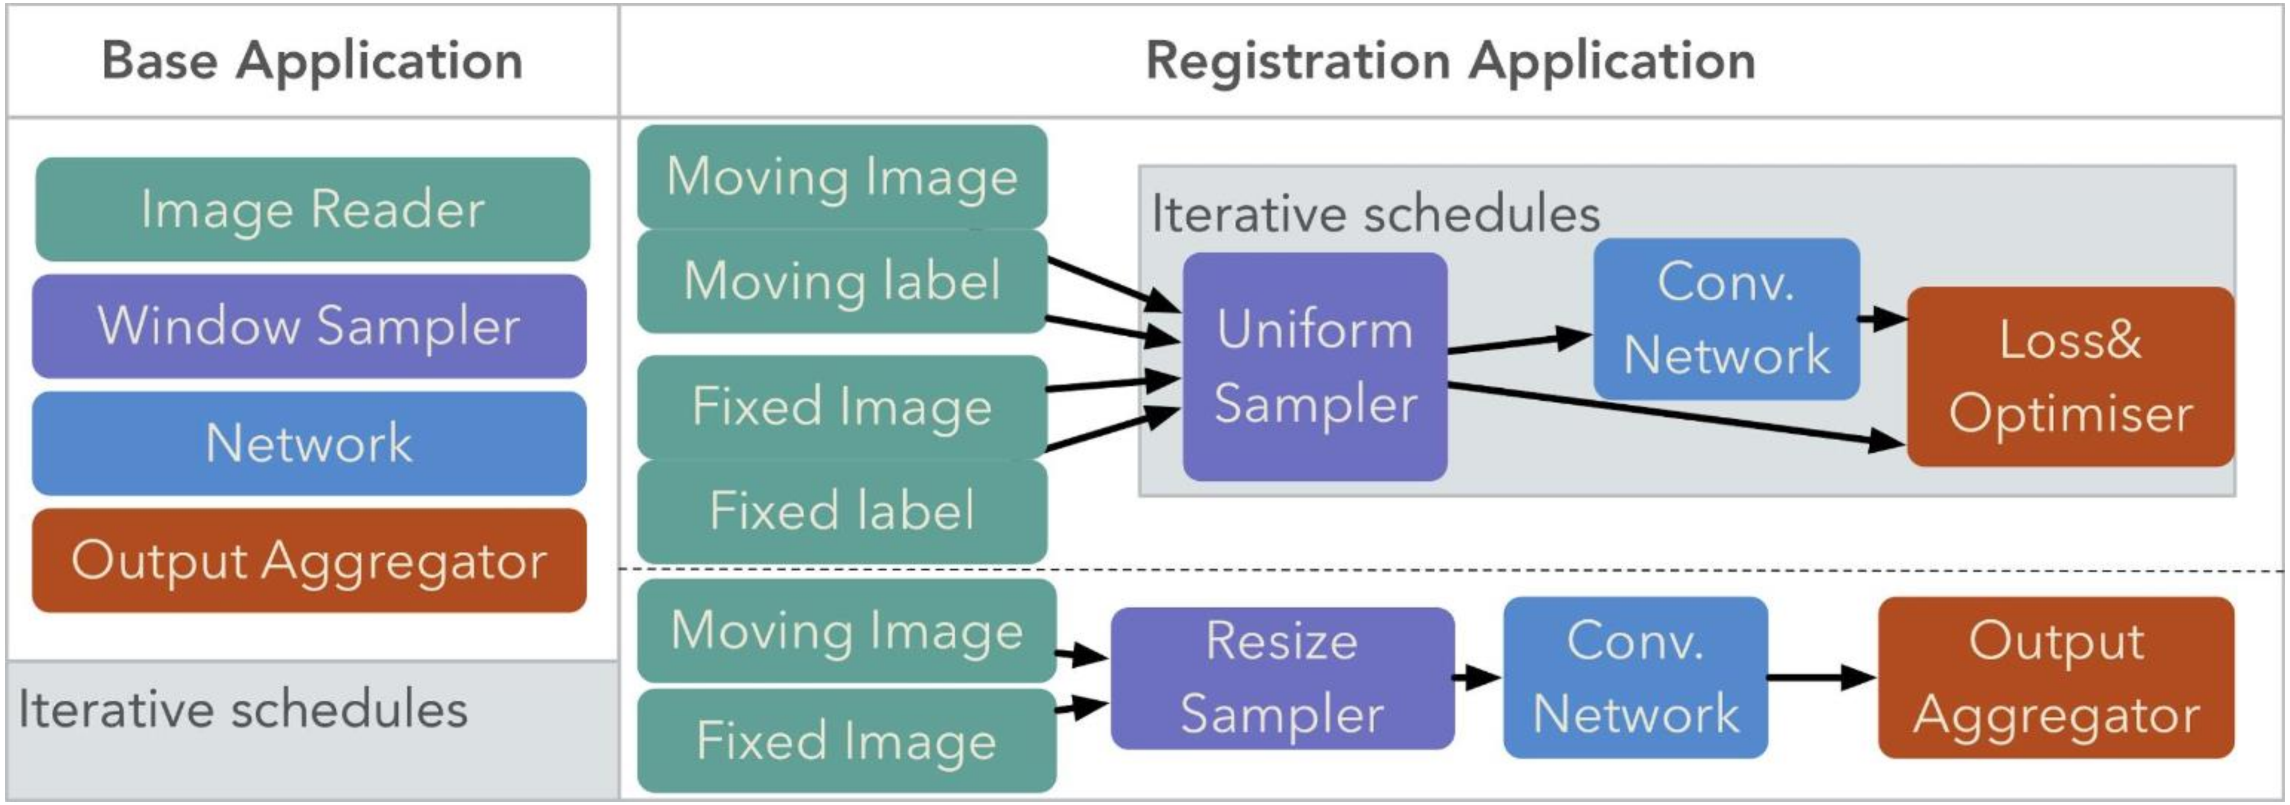

Model zoo

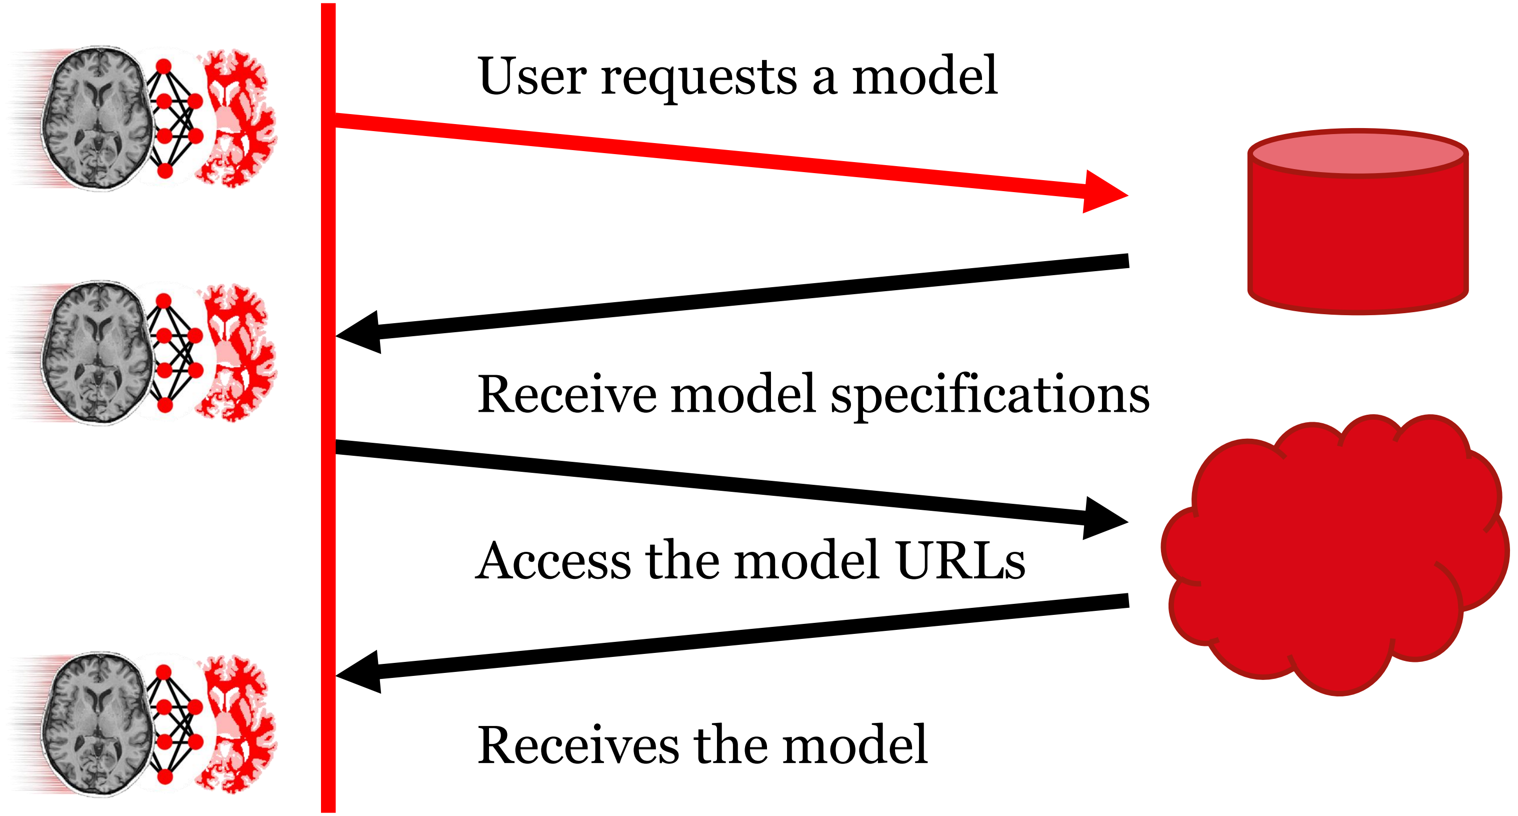

- Sharing networks/apps
  - designs (Python code)
  - trained network weights
  - hyperparameters
  - demo input data
  - Readme file

```
pip install niftynet tensorflow
net_download dense_vnet_abdominal_ct_model_zoo
net_segment inference -c ./dense_vnet_abdominal_ct/config.ini
```

I/O: Image Loader

- Multi image-format loader
  - Uses tf.data API
- Supports multimodal inputs
  - Internally or externally
  - Resolution matching
- Handling a set of image volumes
  - Subject or filename grouping
  - Handling missing modalities
- Preprocessing
  - Handling NIfTI/MHD/DICOM file headers
    - Resampling
    - Reorientation
    - Lazy Sampling
  - Intensity normalisation
    - Physics/Model based

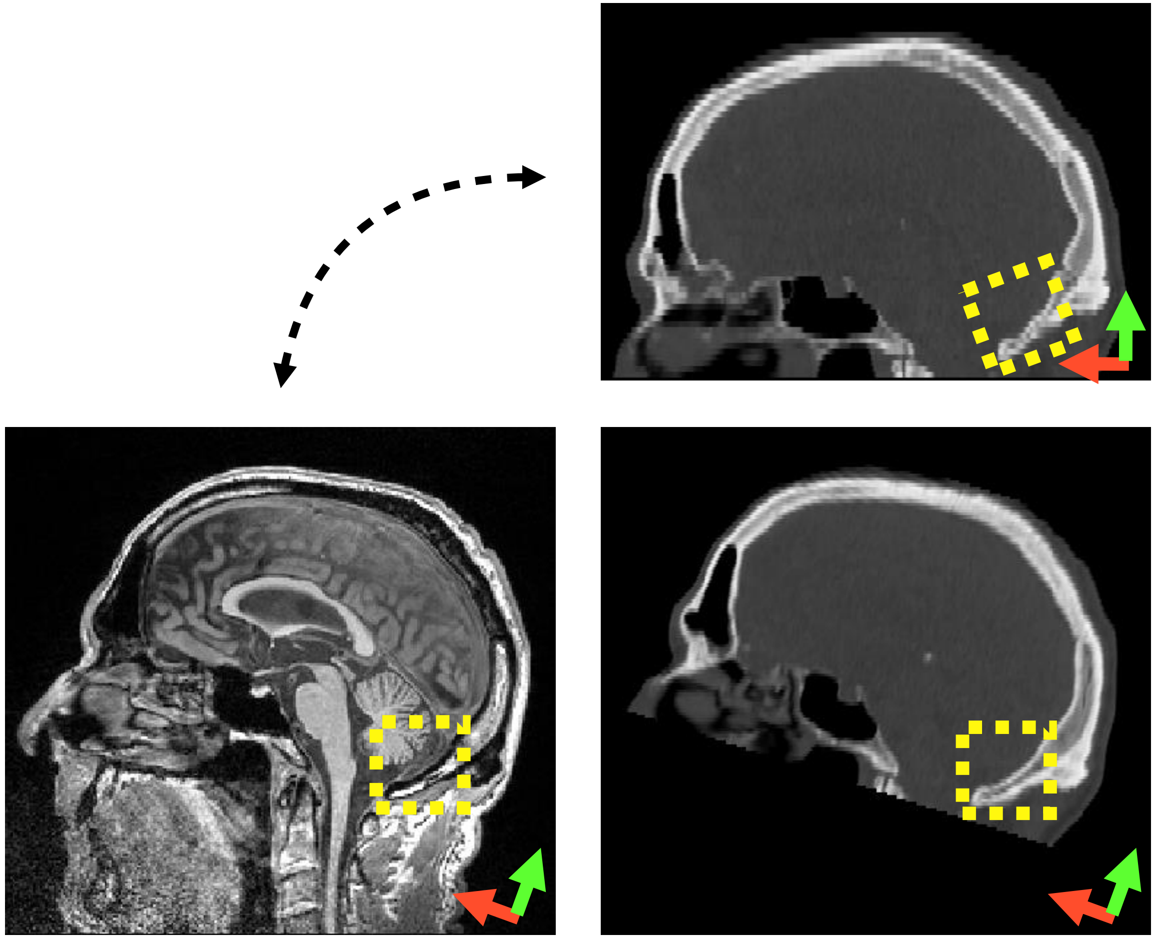

I/O: Window sampling and output aggregation

- Window properties
  - Size in “voxels” or “mm”
  - Augmentation by composition
- Sampling
  - Uniform
  - Label Constrained
    - Sample only from areas with specific labels
    - Prescribe certain label sampling rates
  - Frequency Sampler
    - Sample a location given an externally defined map
    - Sample from locations with large errors
- Aggregation
  - Uniform & Overlapping (Effective Receptive Field)
  - Uncertainty Sampling

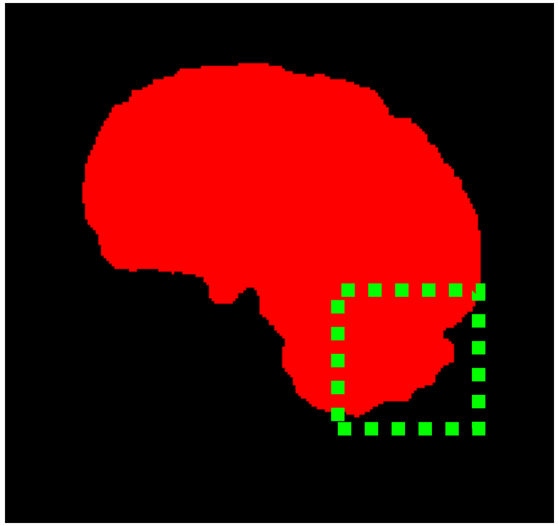

I/O: Data augmentation

- Training with data augmentation
  - Application-dependent
- Geometrical augmentation
  - Rotation, Translation, Mirror
  - Random elastic deformation
  - Biologically-inspired elastic deformation
- Intensity augmentation
  - Histogram/Physics
  - Noise
  - Point-spread-function
  - Artefacts
  - Pathology/lesions

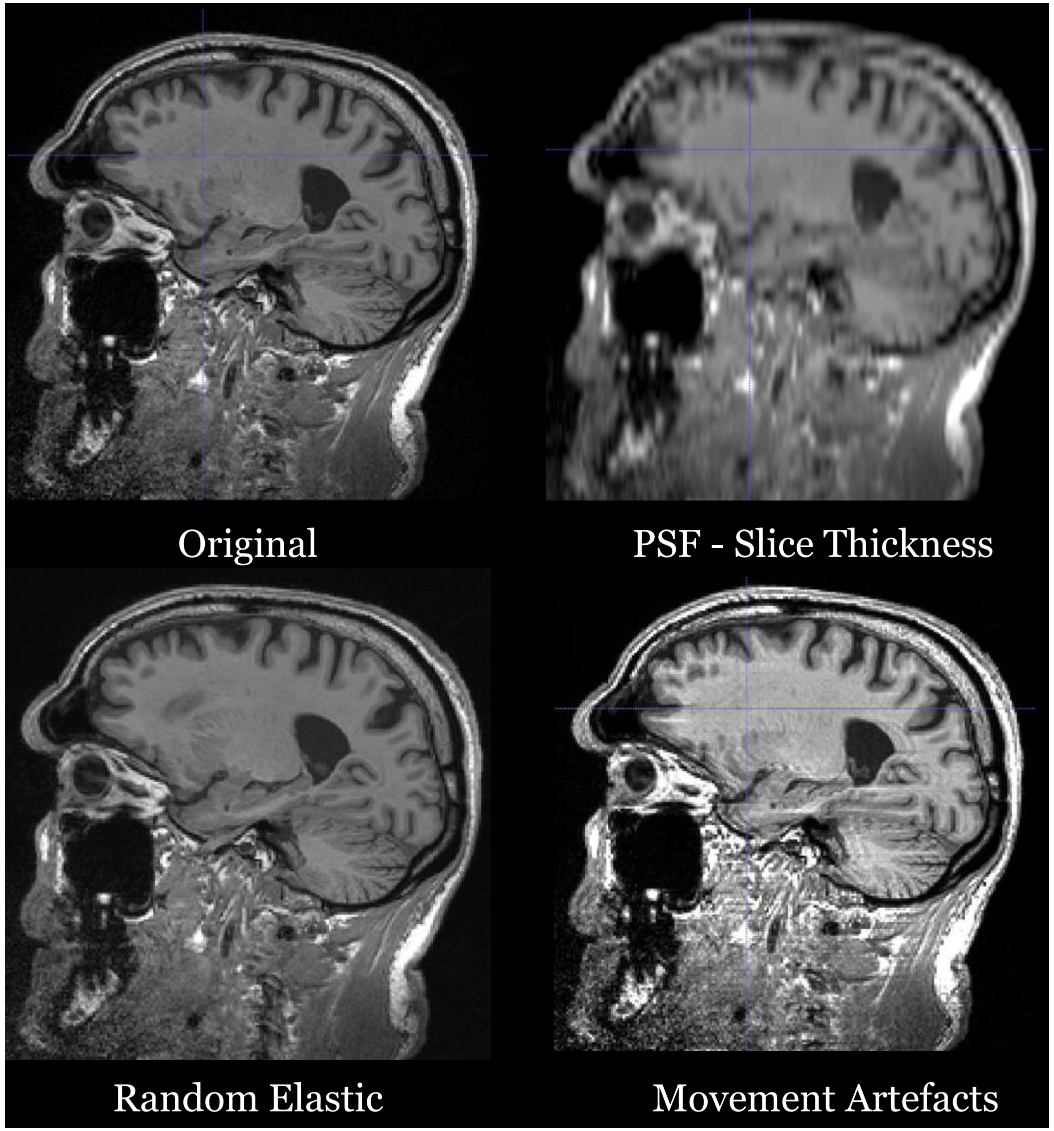

Network: Designing loss functions

- Losses:
  - Categorical
    - Cross-Entropy
    - Dice (Standard, Generalised, Wasserstein)
    - Sensitivity/Specificity
  - Continuous
    - L2/L1
    - Huber
  - Adversarial
  - Variational
    - KLD
- Loss Types
  - Image-wide (e.g. Classification/GAN)
  - Voxel-wise (e.g. Segmentation/Synthesis)
  - Weighted & probabilistic variants

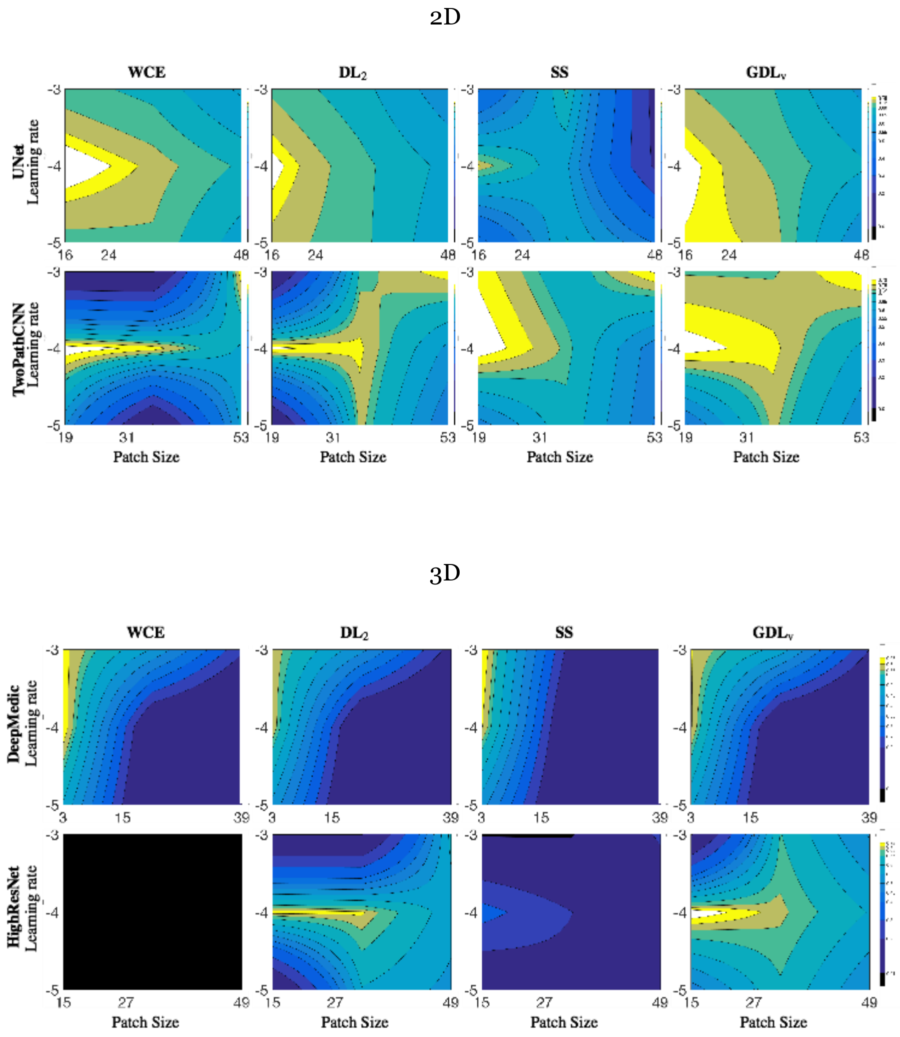

Network: Architectures

- Image-to-image: 2D, 3D, 4D (multimodal)
  - UNet
  - VNet
  - Highway Network
  - DeepMedic
  - HighResNet
- Generative/AutoEncoders
  - AE, dAE, VAE
  - GAN
- Image-to-label
- Multi-task Learning

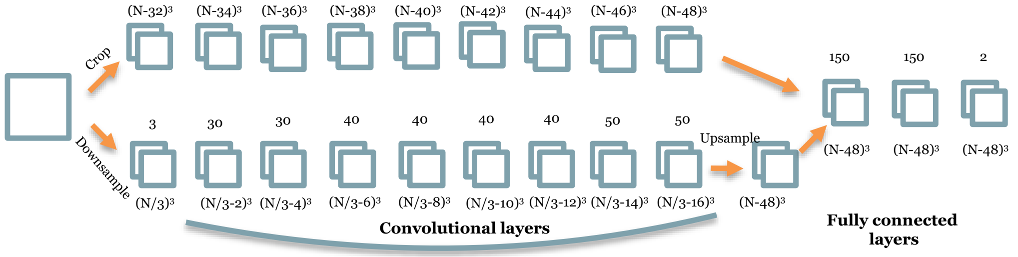

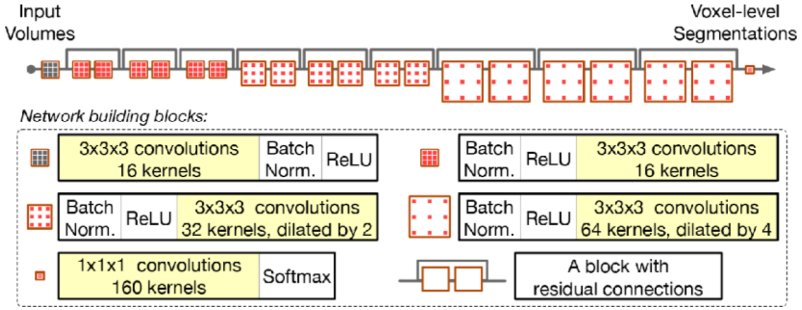

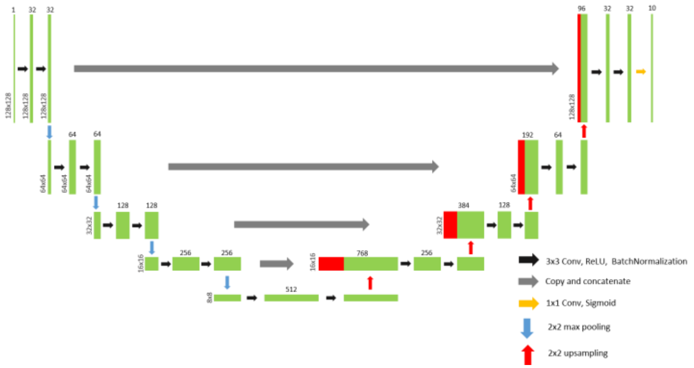

Evaluation: Standardised and validated

- Tensorboard Integration
- Image-level
  - Categorical
    - Overlap: Dice/Jaccard
    - Distance: Hausdorf/MSD
    - Statistical: Sensitivity/Specificity/Recovery
  - Continuous
    - Direct: Mean Absolute Error/L2
    - Perceptual: PSNR/SSIM
- Object Level (Categorical)
  - Volume: Size
  - Overlap, Distance and Statistical metrics
  - F1 stats
- Pixel-level
  - Generation of error maps
- Hyperparameter Optimisation
  - Grid, Random and Divide-and-Conquer Search

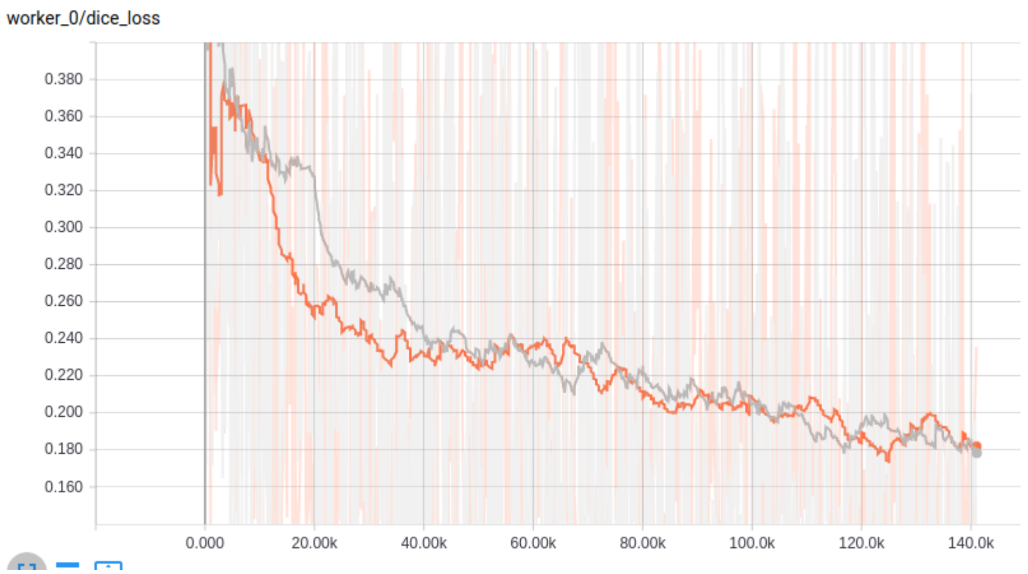

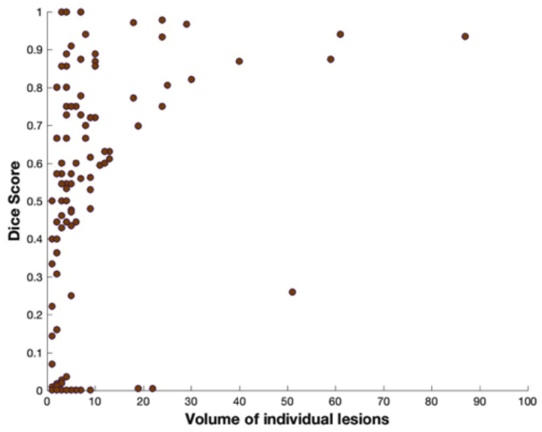

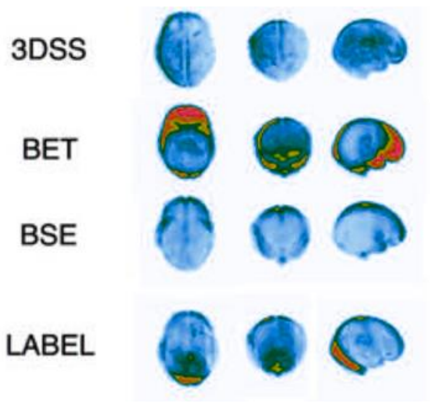

What next?

Checkout the repo.
Run model zoo entries

Modify config.
Run with your data

New networks
Contributing to model zoo

New modules/tasks
Send us pull requests

Thank you

NiftyNt: An open-source convolutional neural networks platform for research in medical image analysis and image-guided therapy

Presented by Tom Vercauteren

http://niftynet.io
https://github.com/NifTK/NiftyNet