Imperial College London

BioMedIA

DLTK

Deep Learning Toolkit for Medical Imaging

Martin Rajchl

Imperial College London

BioMedIA

open science, open source

dltk.github.io

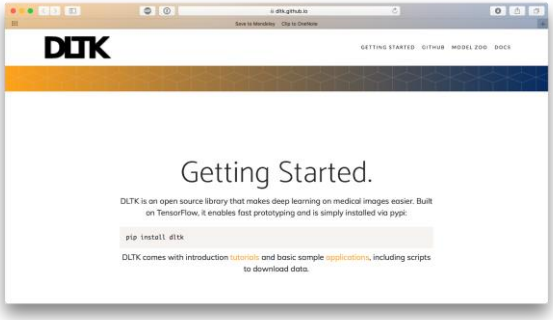

github.com/dltk/dltk

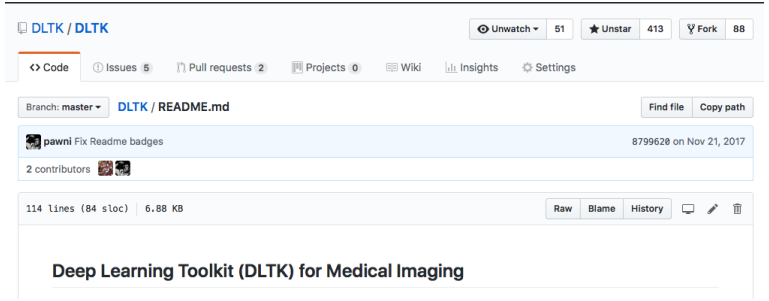

dltk.github.io/DLTK/

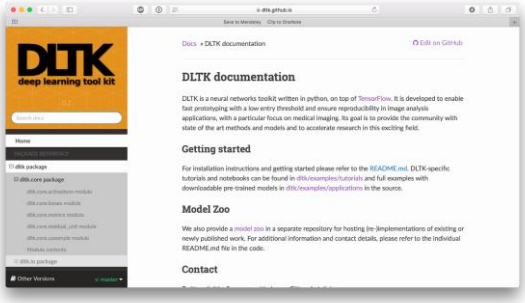

github.com/dltk/models

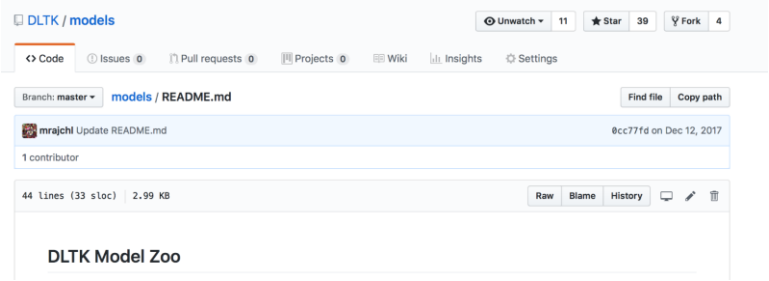

scope

comparison

DLTK

- hackable
- low-threshold access
- speciality ops for MI
- scalable
- tensorflow

The DeepMedic

- out-of-the-box
- image segmentation
- theano

NiftyNet

- full applications
- high-level config
- modules
- tensorflow

Caffe

Caffe2

Chainer

Cognitive Toolkit

MATLAB

mxnet

PaddlePaddle

PYTORCH

TensorFlow

...

features

Imperial College London

BioMedIA

introduction & tensorflow blog

Writer

Edit 74

744

DLTK

deep learning tool kit

website: https://dltk.github.io; source: https://github.com/DLTK/DLTK;

Overview

What is biomedical image analysis and why is it needed? Biomedical images are measurements of the human body on different scales (i.e. microscopic, macroscopic, etc.). They come in a wide variety of imaging modalities (e.g. a CT scanner, an ultrasound machine, etc.) and measure a physical property of the human body (e.g. radiodensity, the opacity to X-rays). These images are interpreted by domain experts (e.g. a radiologist) for clinical tasks (e.g. a diagnosis) and have a large impact on decision making of physicians.

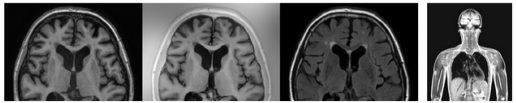

744 3

Next story

Imperial College London

BioMedIA

tutorials & low entry threshold

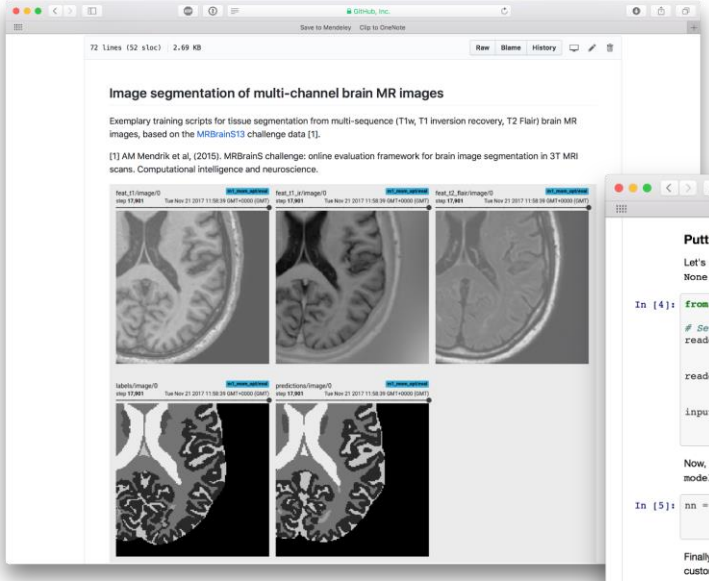

Image segmentation of multi-channel brain MR images

Exemplary training scripts for tissue segmentation from multi-sequence (T1w, T1 inversion recovery, T2 Flair) brain MR images, based on the MRBrainS13 challenge data [1].

[1] AM Mendrik et al, (2015). MRBrainS challenge: online evaluation framework for brain image segmentation in 3T MRI scans. Computational intelligence and neuroscience.

https://github.com/DLTK/DLTK/tree/master/examples/applications/MRBrainS13_tissue_segmentation

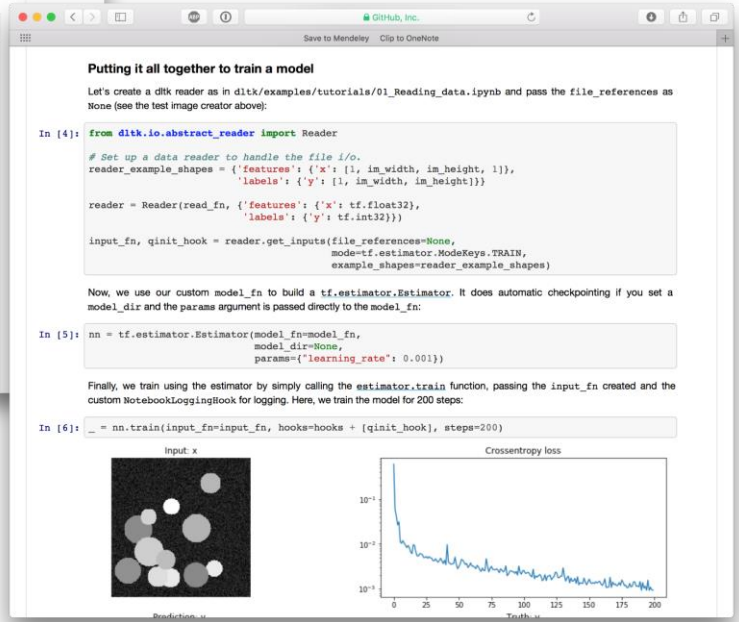

Putting it all together to train a model

Let's create a dltk reader as in `dltk/examples/tutorials/01_Reading_data.ipynb` and pass the `file_references` as `None` (see the test image creator above):

```
In [4]: from dltk.io.abstract_reader import Reader

        # Set up a data reader to handle the file i/o.
        reader_example_shapes = {'features': {'x': [1, im_width, im_height, 1]},
                                 'labels': {'y': [1, im_width, im_height]}}

        reader = Reader(read_fn, {'features': {'x': tf.float32},
                                  'labels': {'y': tf.int32}})

        input_fn, qinit_hook = reader.get_inputs(file_references=None,
                                                 mode=tf.estimator.ModeKeys.TRAIN,
                                                 example_shapes=reader_example_shapes)
```

Now, we use our custom `model_fn` to build a `tf.estimator.Estimator`. It does automatic checkpointing if you set a `model_dir` and the `params` argument is passed directly to the `model_fn`:

```
In [5]: nn = tf.estimator.Estimator(model_fn=model_fn,
                                    model_dir=None,
                                    params={"learning_rate": 0.001})
```

Finally, we train using the estimator by simply calling the `estimator.train` function, passing the `input_fn` created and the custom `NotebookLoggingHook` for logging. Here, we train the model for 200 steps:

```
In [6]: _ = nn.train(input_fn=input_fn, hooks=hooks + [qinit_hook], steps=200)
```

Input: x

Crossentropy loss

https://github.com/DLTK/DLTK/blob/master/examples/tutorials/02_Building_a_model_fn.ipynb

Imperial College London

BioMedIA

data download scripts

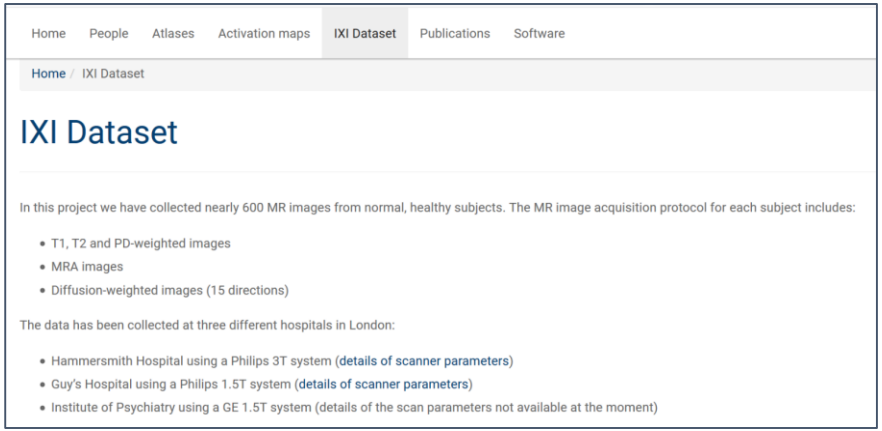

http://brain-development.org/ixi-dataset/

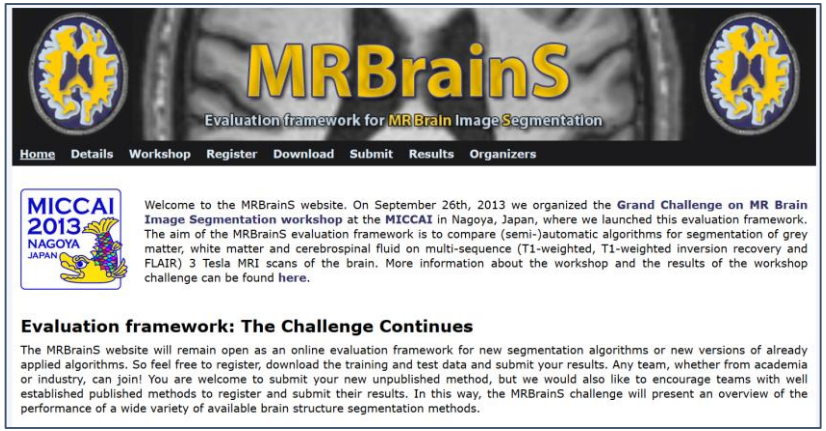

http://mrbrains13.isi.uu.nl/

1. Download
2. Extraction
3. Parse info, labels, etc.
4. Validate
5. Resampling

- Register
- Download
- Copy to data/
- Run example

Imperial College London

BioMedIA

full example applications

Generative modeling of 3T T1 brain MR images

Exemplary training script for generative modeling of T1w brain MR images, based on the IXI dataset [1]. We are using a custom DCGAN-like [2] architecture.

[1] IXI – Information eXtraction from Images (EPSRC GR/S21533/02)
[2] Radford et al. Unsupervised representation learning with deep convolutional generative adversarial networks. 2015

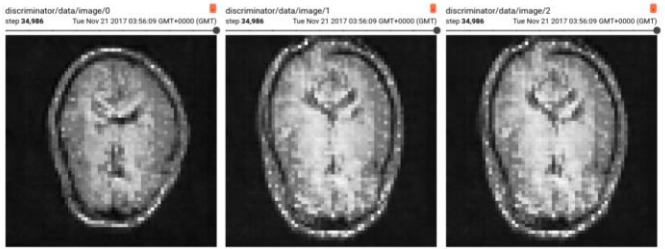

Simple image super-resolution on T1w brain MR images

Exemplary training scripts for image super-resolution from T1w brain MR images, based on the IXI dataset [1].

[1] IXI – Information eXtraction from Images (EPSRC GR/S21533/02)

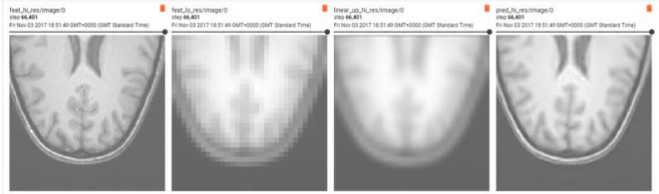

Image segmentation of multi-channel brain MR images

Exemplary training scripts for tissue segmentation from multi-sequence (T1w, T1 inversion recovery, T2 Flair) brain MR images, based on the MRBrainS13 challenge data [1].

[1] AM Mendrik et al. (2015). MRBrainS challenge: online evaluation framework for brain image segmentation in 3T MRI scans. Computational intelligence and neuroscience.

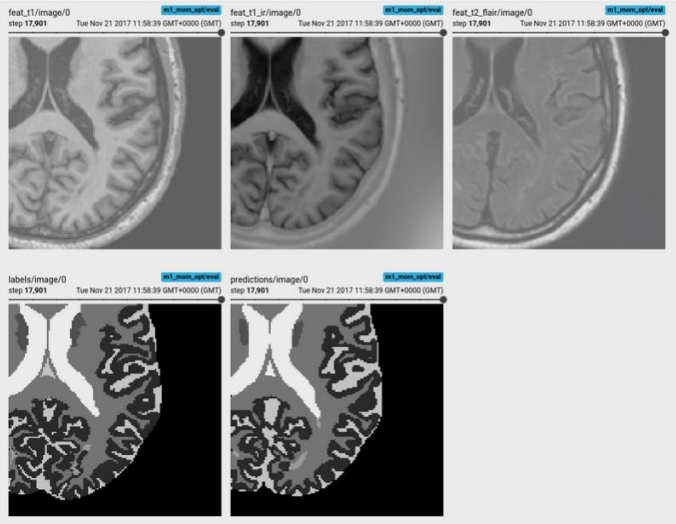

Age regression from 3T T1w brain MR images

Exemplary training and evaluation scripts for regression from T1w brain MR images, based on the IXI dataset [1].

[1] IXI – Information eXtraction from Images (EPSRC GR/S21533/02)

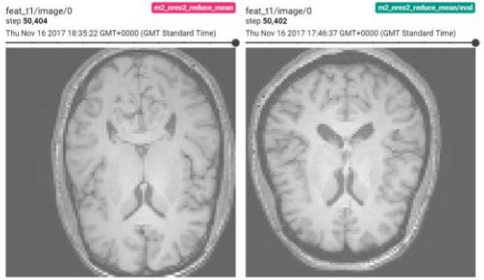

Imperial College London

BioMedIA

high-level training & deploy scripts

```
print('Starting training...')
try:
    for _ in range(MAX_STEPS // EVAL_EVERY_N_STEPS):
        nn.train(
            input_fn=train_input_fn,
            hooks=[train_qinit_hook, step_cnt_hook],
            steps=EVAL_EVERY_N_STEPS)

        if args.run_validation:
            results_val = nn.evaluate(
                input_fn=val_input_fn,
                hooks=[val_qinit_hook, val_summary_hook],
                steps=EVAL_STEPS)
            print('Step = {}; val loss = {:.5f};'.format(
                results_val['global_step'],
                results_val['loss']))
```

High-level tf.estimator API for training

```
cat deploy.py | wc -l
129
```

Easy to read deploy scripts

```
def model_fn(features, labels, mode, params):
    """Model function to construct a tf.estimator.EstimatorSpec. It creates a
        network given input features (e.g. from a dltk.io.abstract_reader) and
        training targets (labels). Further, loss, optimiser, evaluation ops and
        custom tensorboard summary ops can be added. For additional information,
         please refer to https://www.tensorflow.org/api_docs/python/tf/estimator/Estimator#model_fn.


    # 5. Return EstimatorSpec object
    return tf.estimator.EstimatorSpec(mode=mode,
                                      predictions=net_output_ops,
                                      loss=loss,
                                      train_op=train_op,
                                      eval_metric_ops=eval_metric_ops)
```

Standardised tf.estimators interface for feed forward networks

Imperial College London

BioMedIA

scalable reference implementations

```
# 1. create a model and its outputs
net_output_ops = residual_unet_3d(
    inputs=features['x'],
    num_classes=NUM_CLASSES,
    num_res_units=2,
    filters=(16, 32, 64, 128),
    strides=((1, 1, 1), (1, 2, 2), (1, 2, 2), (1, 2, 2)),
    mode=mode,
    kernel_regularizer=tf.contrib.layers.l2_regularizer(1e-4))
```

Scaling pre-built network architectures via parameterisation

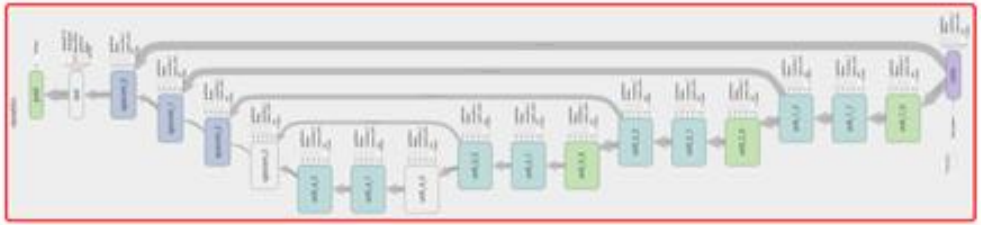

- Pre-built networks
- Validated performance of components
- Independent of input tenor shapes (future: rank)
- Quickly scale to new problem

a quick regression experiment

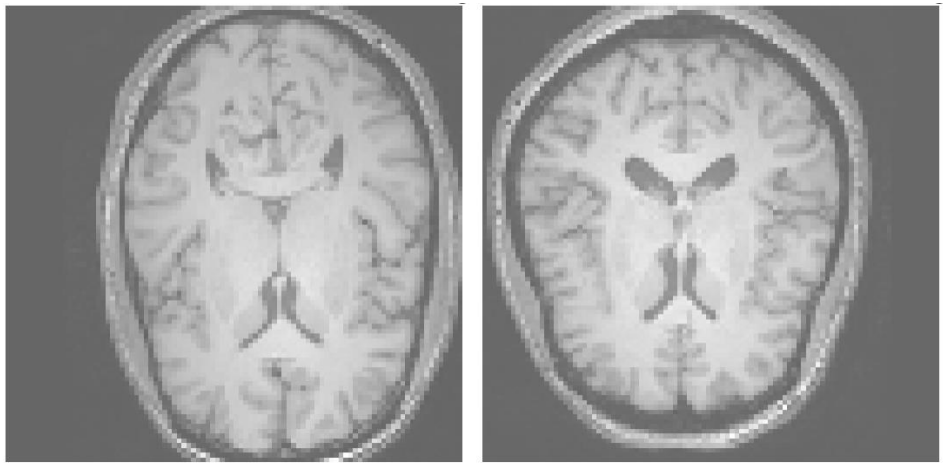

Toy task: Age regression from T1w MR images
IXI database, 2mm isotropic, random crops

IXI dataset, https://brain-development.org/

Imperial College London

BioMedIA

a quick regression experiment

```
{
  "net": "fcn",
  "filters": [16, 64, 128, 256, 512],
  "strides": [[1, 1, 1], [2, 2, 2], [2, 2, 2], [2, 2, 2], [1, 1, 1]],
  "num_residual_units": 3,
  "loss": "balce",
  "learning_rate": 0.001
}
```

JSON configs for experimental setups and high-level parameterisation

Residual Unit

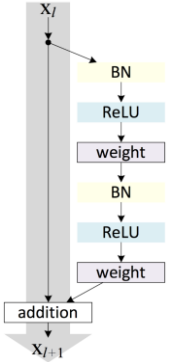

Identity Mappings in Deep Residual Networks, He et al. 2016

Dense Unit

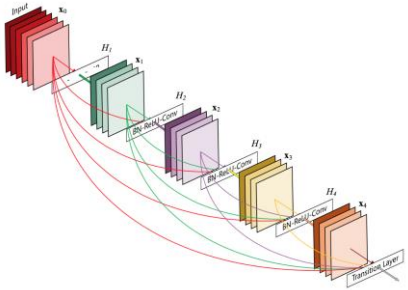

Densely Connected Convolutional Networks, Huang et al. 2016

ResNext Unit

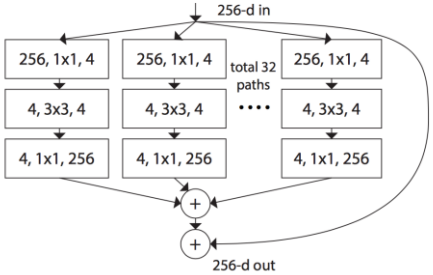

Aggregated Residual Transformations for Deep Neural Networks, Xie et al. 2016

Squeeze-Excitation Unit

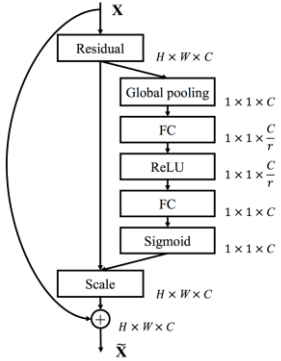

Squeeze-and-Excitation Networks, Hu et al., 2018

a quick regression experiment

	MAE	num weights
Residual	4.929 yrs	9,6m
Dense	5.189 yrs	1,8m
ResNext	4.764 yrs	2,3m
SE-ResNext	4.199 yrs	2.4m

AGE	mean	std	min	max
all	47.35	16.76	20.17	81.94
train	48	17.14	20.17	81.94
test	43.89	14.02	25.53	71.21

validated performance

Kamnitsas, EMMA, BRATS challenge 1st, 2017

Free Competition

ID	name	entity	status	team	user+ID	Dice	Mean+Surface+Distance	Hausdorff+Distance
9631575	CNN Test	syn10513440	SCORED		Nick Pawlowski (pawni)	0.81535	1.8613	62.8724
4650177	IMI_deeds_SSC_jointCL	syn4650176	SCORED	IMI_deedsSSC	Mattias Heinrich (mattias)	0.78963	2.2619	25.5037
9624748	Zhou	syn10205118	SCORED	Google	Steve Damon (damons)	0.77832	2.9252	32.6082
4635823	IMI_deedsSSC_joint	syn4635822	SCORED	IMI_deedsSSC	Mattias Heinrich (mattias)	0.77013	2.4537	29.0027
5592321	deepseg_rb	syn5592320	SCORED	IACV_CHALMERS	Måns Larsson (mns.larsson)	0.75723	3.2975	27.799
5598864	deepseg_rb	syn5592320	SCORED	IACV_CHALMERS	Steve Damon (damons)	0.75723	3.2975	27.799
4732859	CLSIMPLEJLF_organwise	syn4732857	SCORED	MASI	Zhoubing Xu (xuz8)	0.72289	2.9071	32.0399

Pawlowski, synapse challenge, multi-organ segm., 1st place,, 2017

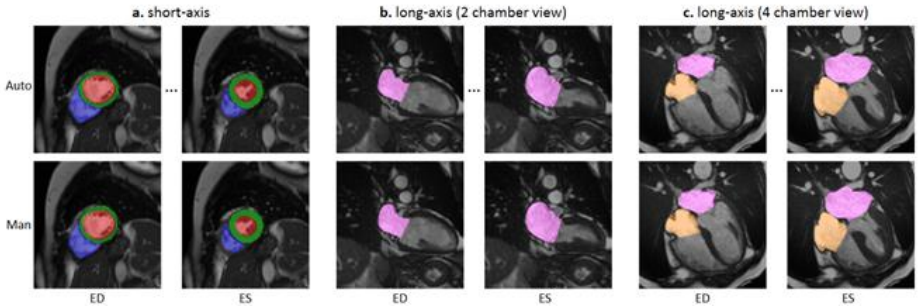

Bai, Human-level performance on UK Biobank, 2018

pre-processing & augmentation

Intensity normalisation (volume stats!):

- -1 to 1 (CT, quantitative images)
- Whitening zero mean, unit std (weighted MRI)
- 0 to 1

Spatial normalisation
(computationally expensive, typically offline)

- Resampling/Reslicing
- Registration (e.g. MNI)

Bias correction

Pre-processing

Add noise:

- Gaussian noise
- Random offsets

Flipping (where meaningful)

Random (elastic) deformations

Random cropping (ideally class-balanced, c.f. sampling)

Augmentation

class balancing

Under-represented classes (e.g. pathologies, small anatomical structures, etc) could either be undersampled (e.g. random sampling) or under-penalised (e.g. mean pixel-wise losses).

loss-based

- median-frequency reweighted
- batch stats reweighted
- mean smooth Dice/MIOU loss

sampling-based

- fixed count sampling
- moving stats sampling

For a performance comparison, see Pawlowski et al. NIPS WS 2017.

model zoo

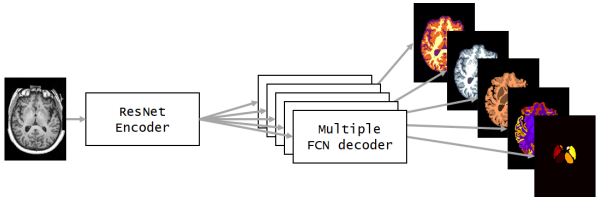

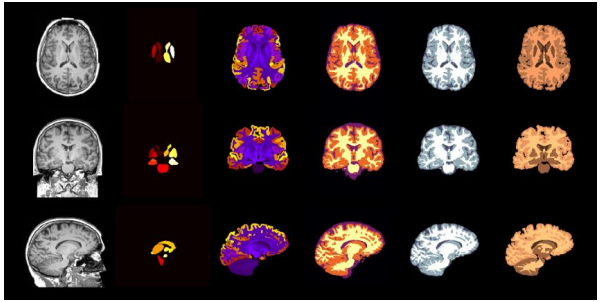

NeuroNet, Rajchl, 2018

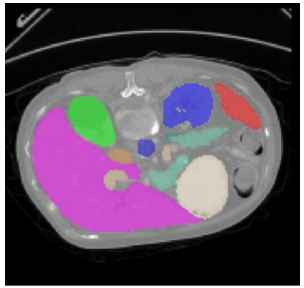

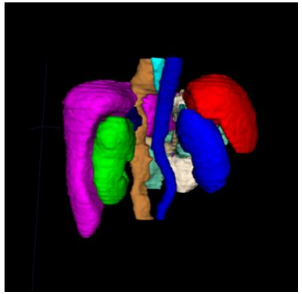

Pawlowski, 2017

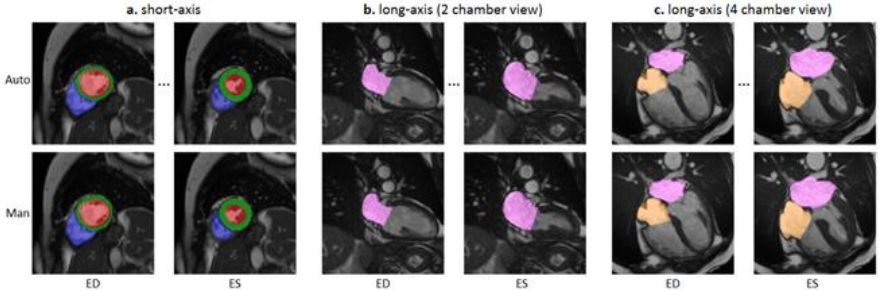

Bai, 2018

Rajchl, 2016

design choices

Imperial College London

BioMedIA

seamless integration into Tensorflow

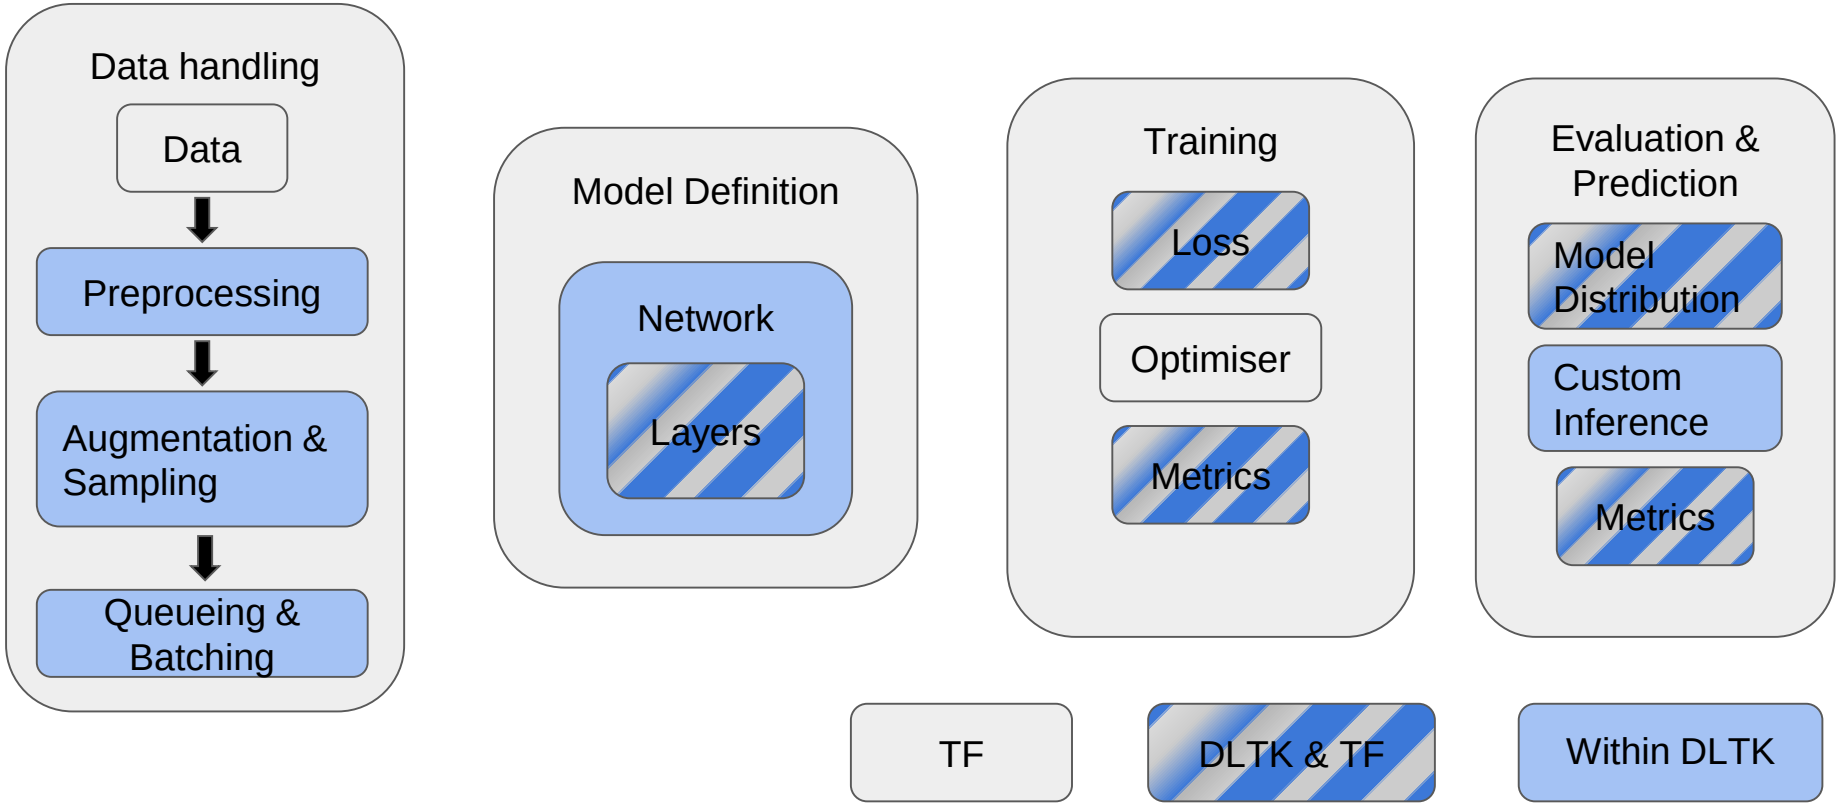

data i/o

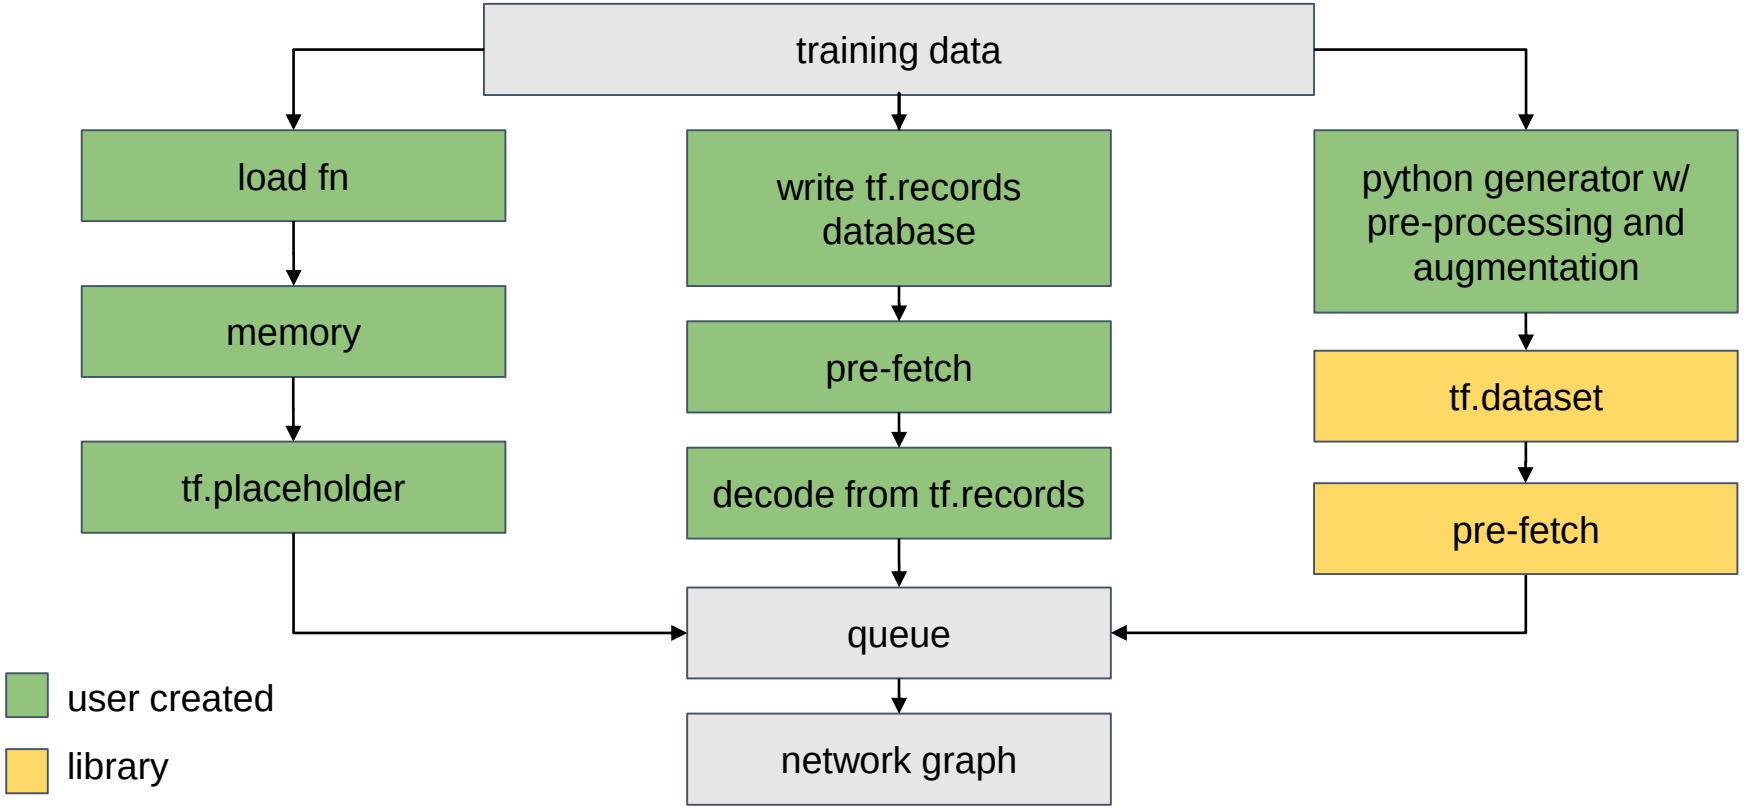
training data
load fn
memory
tf.placeholder
write tf.records database
pre-fetch
decode from tf.records
python generator w/ pre-processing and augmentation
tf.dataset
pre-fetch
queue
network graph
user created
library

dltk reader

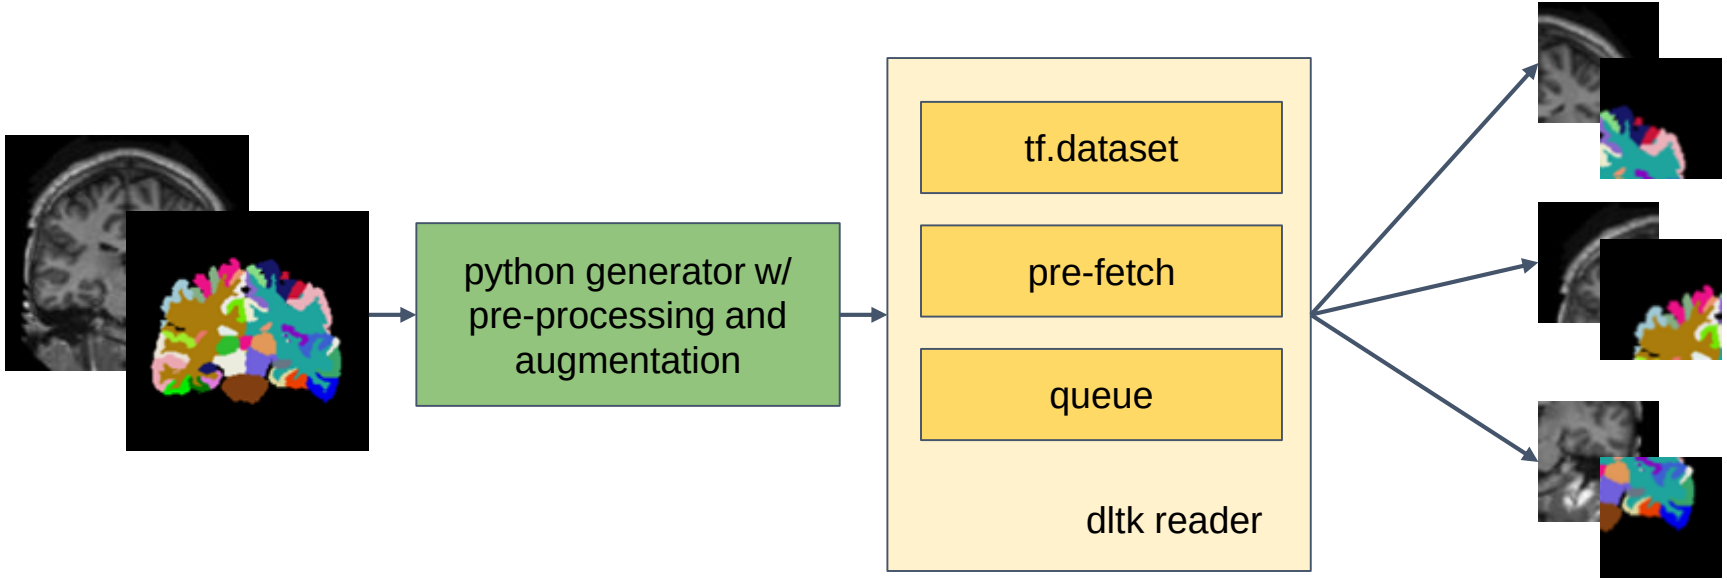

training data

read fn

reader

training examples

exemplary research with DLTK

Imperial College London

BioMedIA

neuronet: fast and robust reproduction of multiple brain image segmentation pipelines

Rajchl et al. MIDL 2018

neuronet

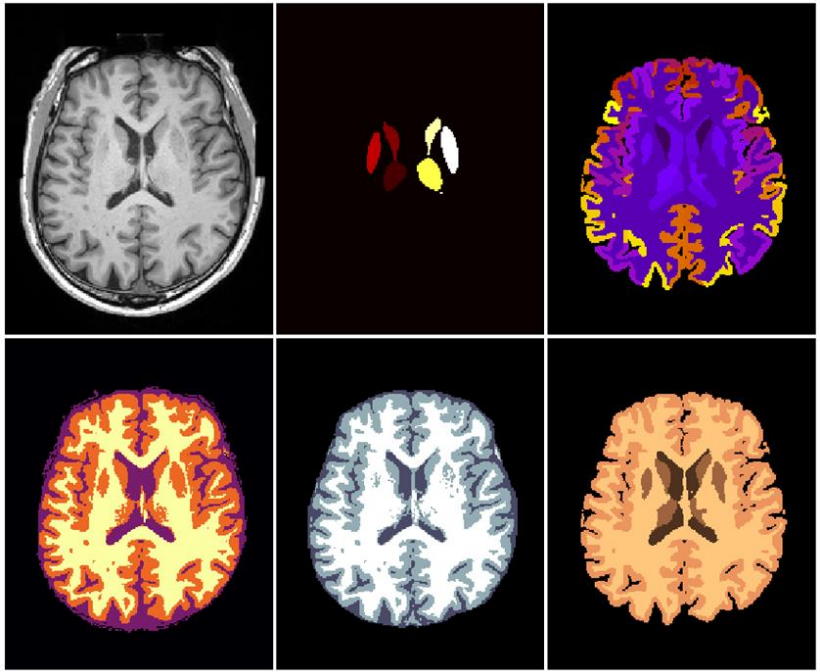

UK Biobank case: Raw T1w MR image, FSL First, MALP-EM, SPM tissue, FSL fast, MALP-EM tissue segmentations.

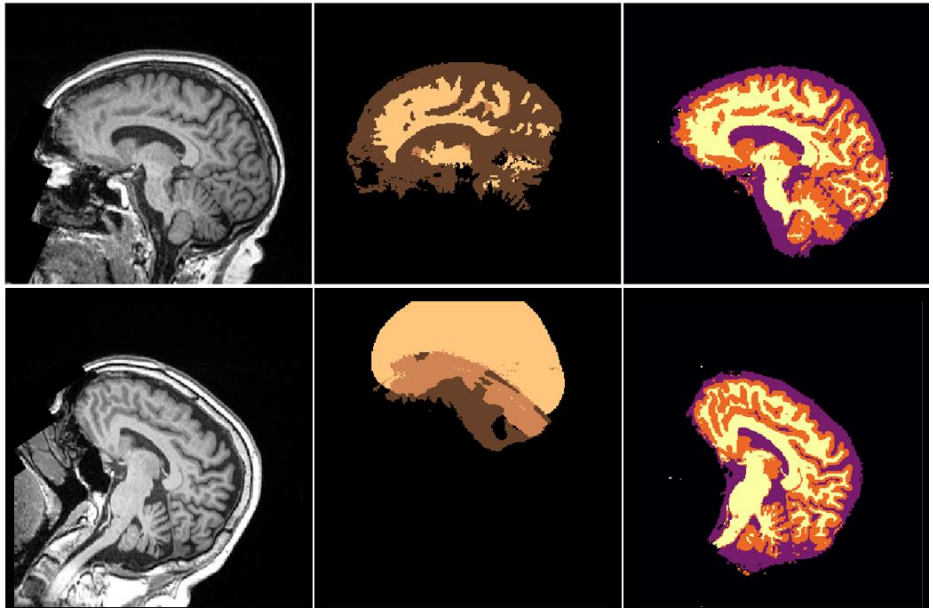

Exemplary failure cases due to head rotation (mid), CNN (right)

Rajchl, MIDL, 2018

neuronet

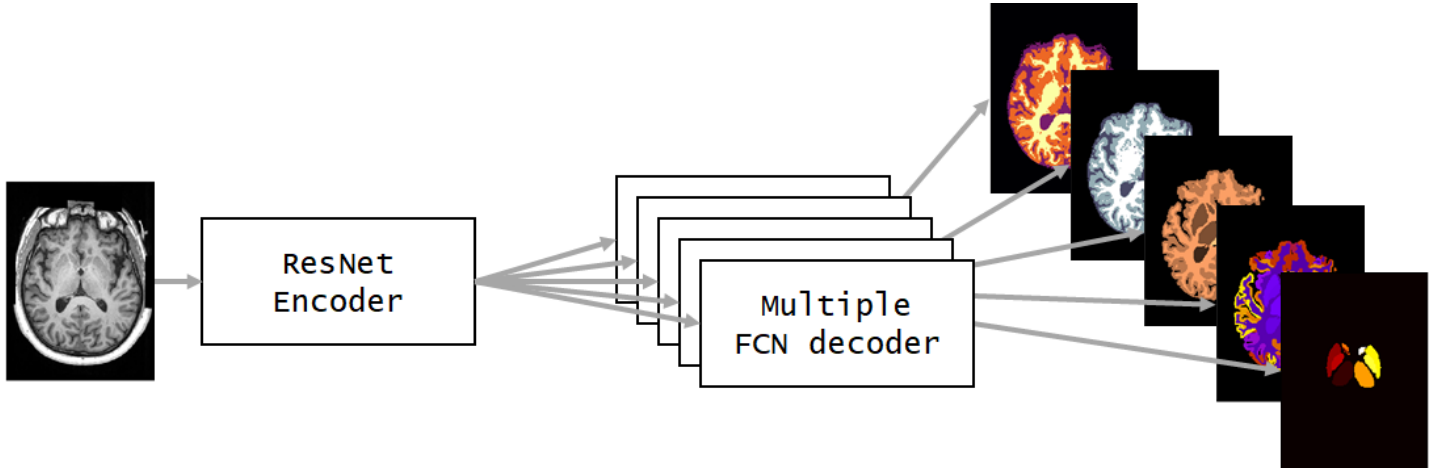

Rajchl, MIDL 2018

Imperial College London

BioMedIA

neuronet

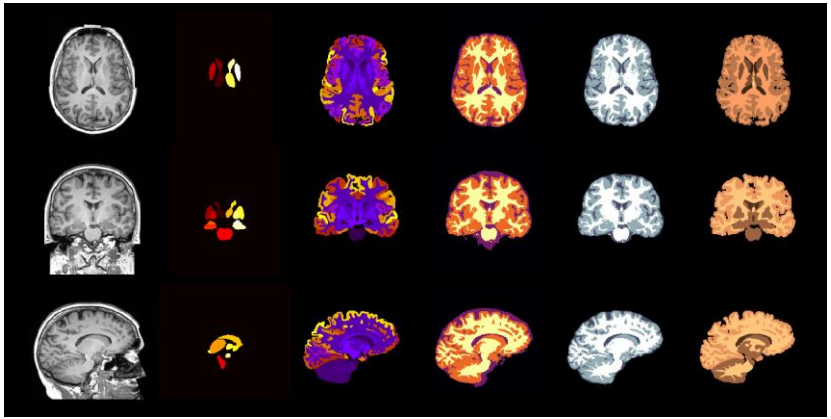

Random UK Biobank case: Raw T1w MR image, FSL First, MALP-EM, SPM tissue, FSL fast, MALP-EM tissue segmentations.

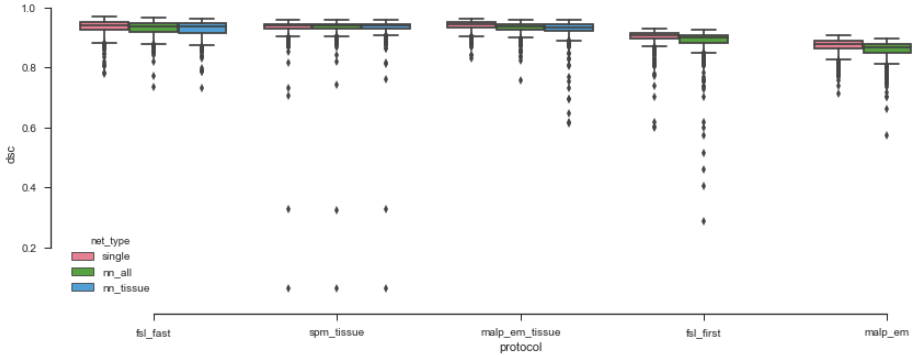

Comparative DSC accuracy for single networks, nn_all and nn_tissue

Rajchl, MIDL, 2018

Imperial College London

BioMedIA

Multi-Modal Learning from Unpaired Images: Application to Multi-Organ Segmentation in CT and MRI

Valindria et al. IEEE WACV 2018

learning from unpaired images

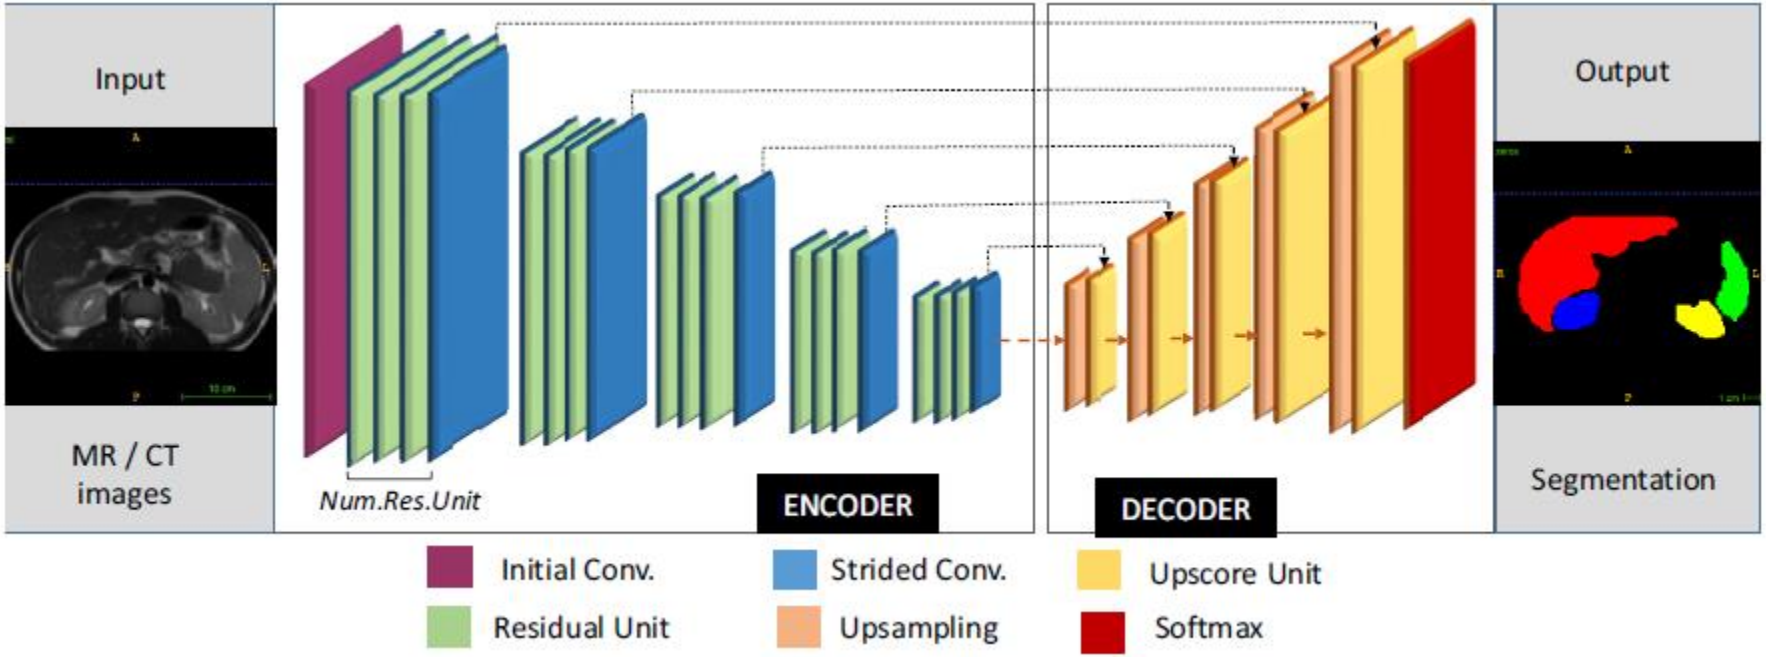

Valindria et al., IEEE WACV, 2018

learning from unpaired images

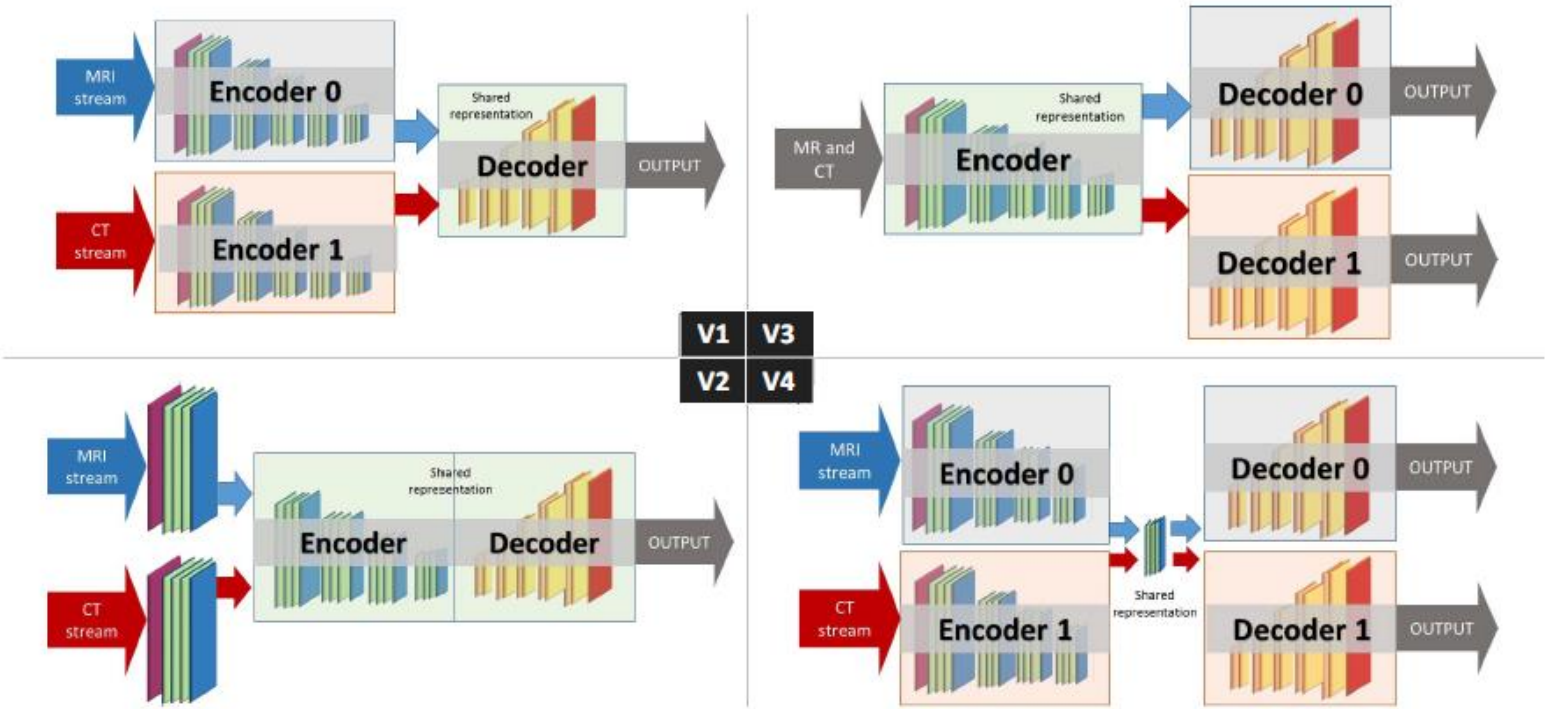

Valindria et al., IEEE WACV, 2018

Imperial College London

BioMedIA

learning from unpaired images

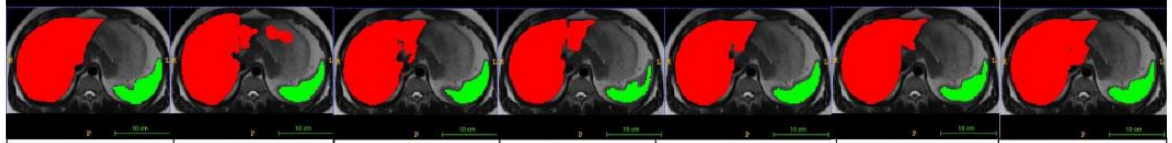

	Ground truth	Individual MR	CTMR-joint	Dualstream V.1	Dualstream V.2	Dualstream V.3	Dualstream V.4
Liver	0.905	0.906	0.902	0.892	0.915	**0.926**	
Spleen	0.856	0.864	0.845	0.852	0.866	**0.893**	

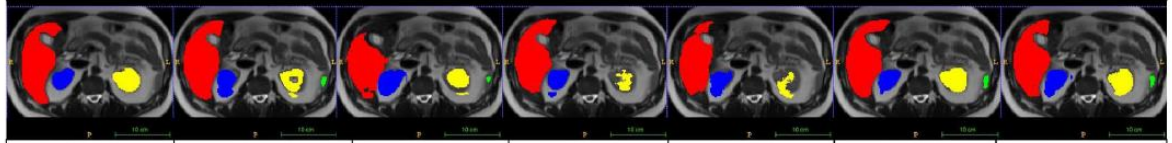

	Ground truth	Individual MR	CTMR-joint	Dualstream V.1	Dualstream V.2	Dualstream V.3	Dualstream V.4
Right Kidney	0.584	0.338	0.568	0.326	0.664	**0.762**	
Left Kidney	0.713	0.329	0.481	0.383	0.758	**0.761**	

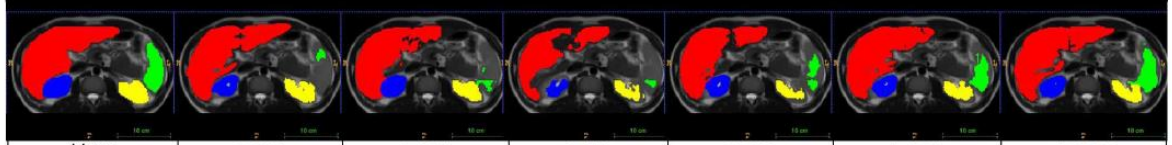

	Ground truth	Individual MR	CTMR-joint	Dualstream V.1	Dualstream V.2	Dualstream V.3	Dualstream V.4
Liver	0.899	0.854	0.811	0.836	0.864	**0.907**	
Spleen	0.750	0.734	0.810	0.786	0.704	**0.840**	
Right Kidney	0.729	0.900	0.891	0.887	0.865	**0.921**	
Left Kidney	0.822	0.878	0.908	0.857	0.839	**0.923**	

Valindria et al., IEEE WACV, 2018

implicit weight uncertainty in neural networks

Pawlowski et al. arXiv 2018

implicit weight uncertainty

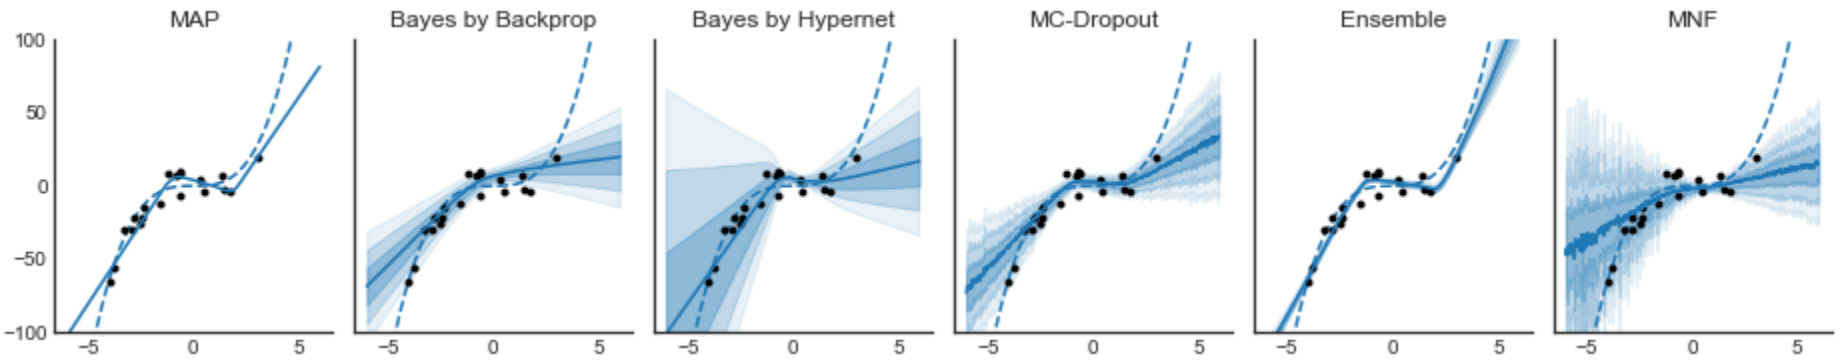

Toy example: model accuracy versus model uncertainty of compared methods

Pawlowski et al, arXiv, 2018

implicit weight uncertainty

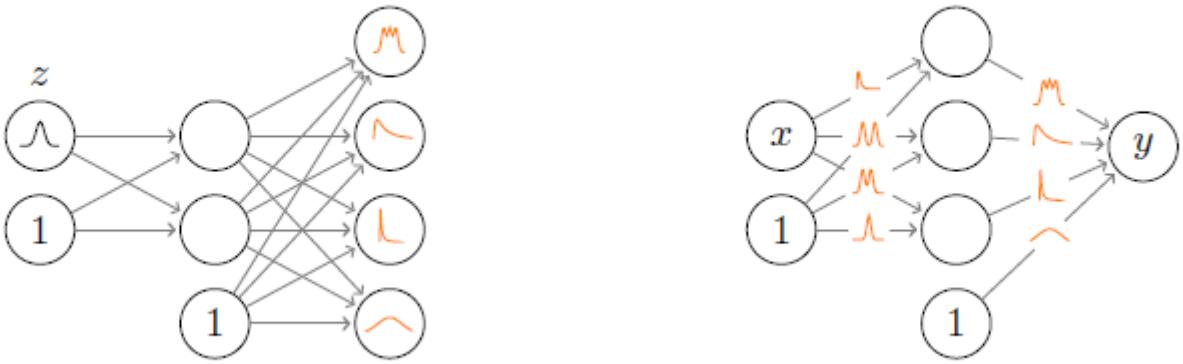

(a) *Hypernetwork G* (b) Main Network

Bayes by Hypernet

Pawlowski et al, arXiv, 2018

Imperial College London

BioMedIA

implicit weight uncertainty

	Error [%]	MNIST AUC	Outlier AUC	Runtime [s]
MAP	0.80	0.99	0.71	710
Ensemble	0.49	0.99	0.65	2721
BbB	0.72	0.97	0.41	2892
Dropout	0.47	0.99	0.58	1224
MNF	0.63	0.99	0.58	21811
BbH (ours)	0.56	0.97	0.48	11504

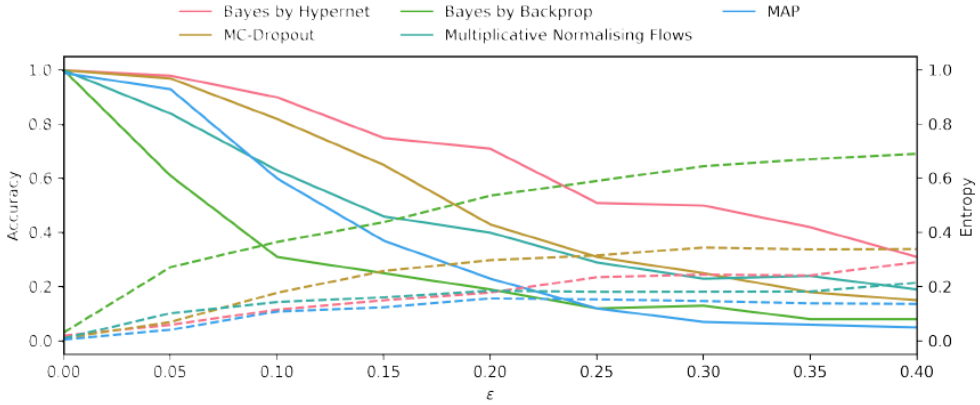

Pawlowski et al, arXiv, 2018

future developments

tf.hub

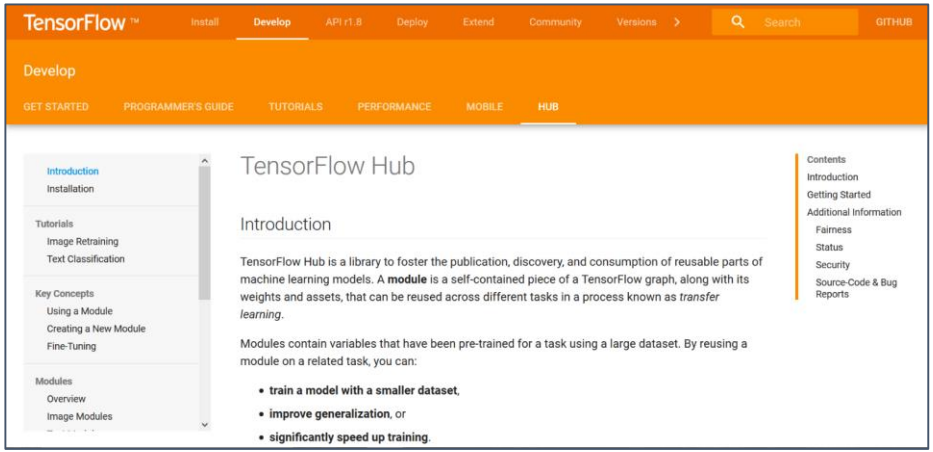


```python
import tensorflow as tf
import tensorflow_hub as hub

with tf.Graph().as_default():
  embed = hub.Module("https://tfhub.dev/google/nnlm-en-dim128-with-normalization/1")
  embeddings = embed(["A long sentence.", "single-word", "http://example.com"])

  with tf.Session() as sess:
    sess.run(tf.global_variables_initializer())
    sess.run(tf.tables_initializer())

    print(sess.run(embeddings))
```

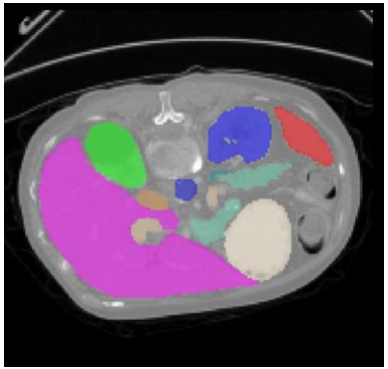

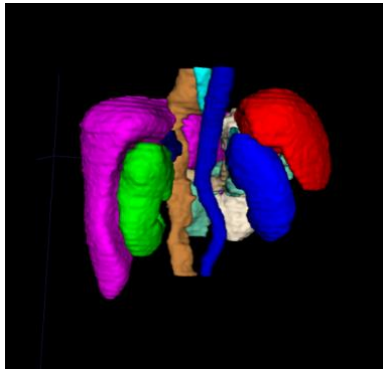

DLTK, Pawlowski, 2017

Conclusion

Try it out!

acknowledgements

Imperial College London

Microsoft Research

erc European Research Council

eHealth

SEVENTH FRAMEWORK PROGRAMME

European Commission Information Society and Media

MRC Medical Research Council

EPSRC Engineering and Physical Sciences Research Council

TensorFlow

NVIDIA

wellcome trust

NHS National Institute for Health Research

upcoming


MEDICAL IMAGING MEETS NIPS
Saturday, December 8th
Palais des Congrès de Montréal, Canada
Call for Abstracts
Up to Three Pages
(excluding refs and acks)
Submissions: 15 October
Notifications: 31 October
Keynote Speakers
Max Welling
Tal Arbel
Olaf Ronneberger
Holger Roth
Organisers
Ender Konukoglu
Hervé Lombaert
Ben Glocker
Marleen de Bruijne
https://sites.google.com/view/med-nips-2018
#mednips


MIDL
London 2019
International Conference on
Medical Imaging with Deep Learning
London, 8 – 10 July 2019
Full papers: 13 Dec 2018
Notification: 28 Feb 2019
Abstracts: April 2019
Notification: May 2019
https://2019.midl.io
Conference Chairs
Tom Vercauteren
Ben Glocker
M. Jorge Cardoso
Program Chairs
Gozde Unal
Ender Konukoglu
Ipek Oguz

Imperial College London

BioMedIA

resources

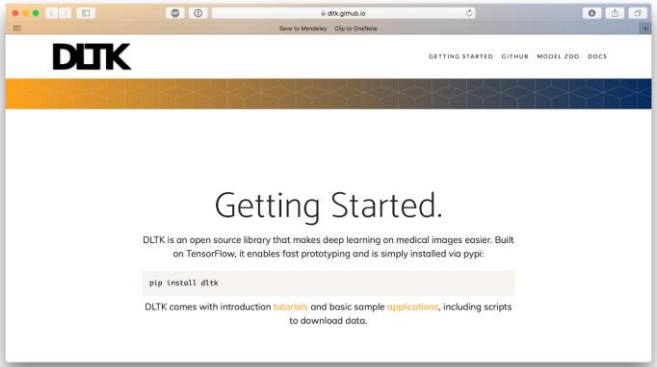

dltk.github.io

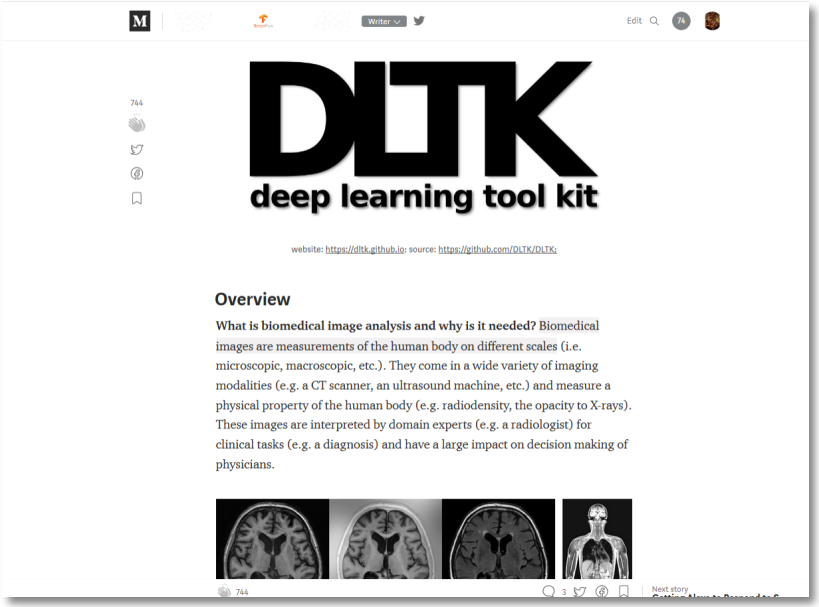

medium.com/tensorflow

Martin Rajchl

m.rajchl@imperial.ac.uk

@m_rajchl

@mrajchl