Why This Educational Course is Not Fully Devoted to Deep Learning

BT Thomas Yeo

Electrical & Computer Engineering & Centre for Sleep and Cognition & Clinical Imaging Research Centre & N.1 Institute for Health

National University of Singapore

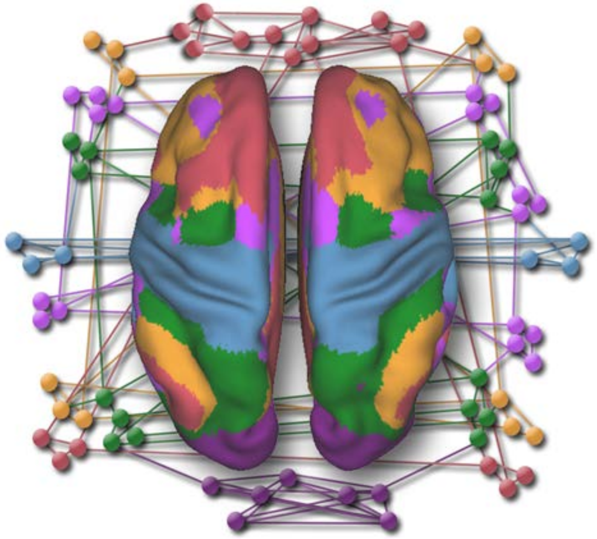

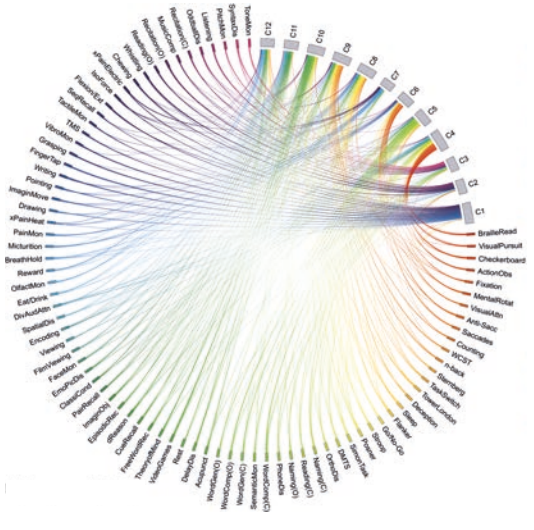

Deep Neural Networks (DNNs) Achieve State-of-the-Art Performance in Many Fields

DeepMind's AlphaGo

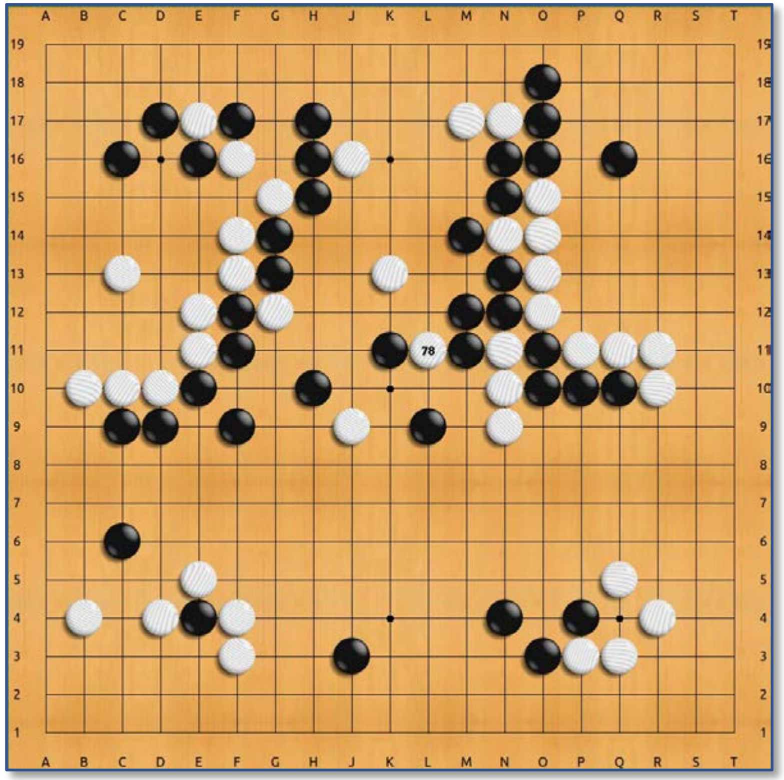

https://en.wikipedia.org/wiki/List_of_Go_terms#Divine_move

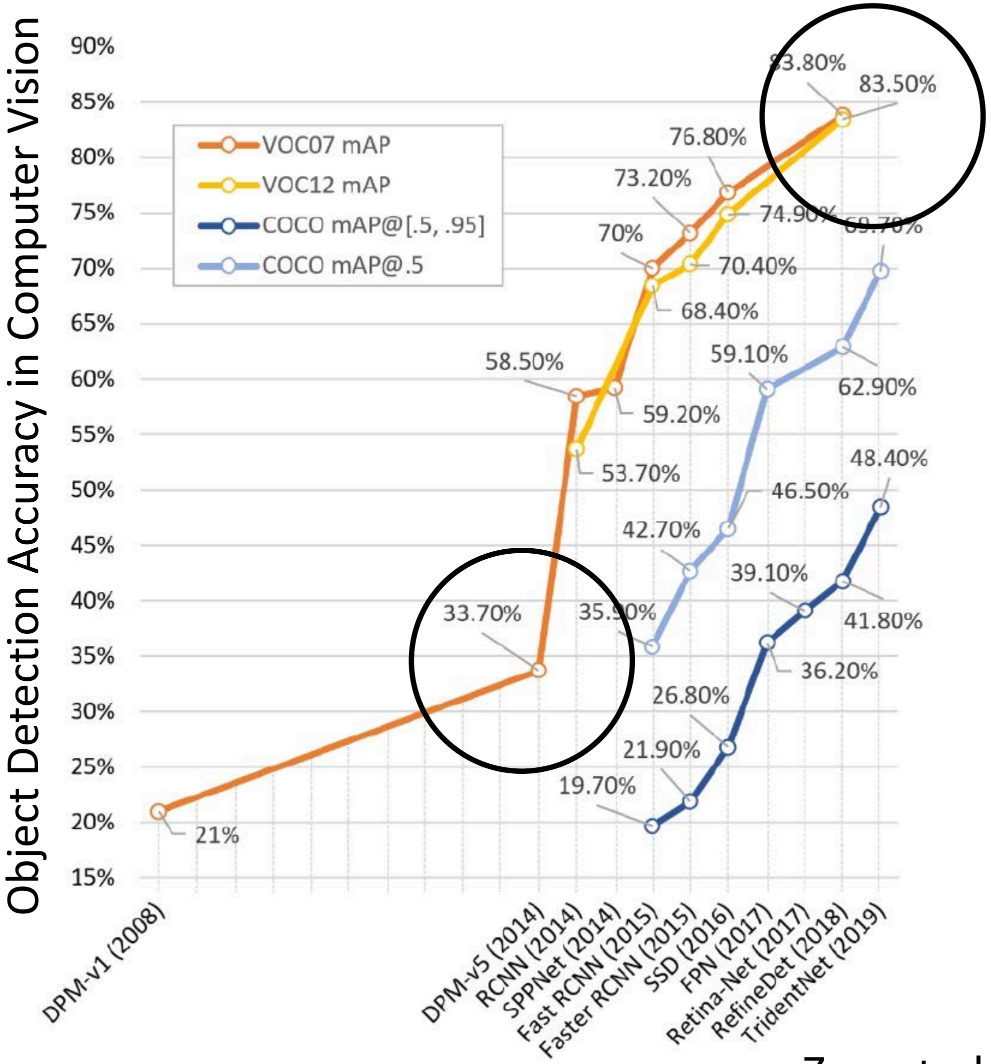

Zou et al., 2019

Classical ML Remains Competitive in Some Domains

Google Research, Scalable & accurate deep learning with electronic health records, NPJ Digital Medicine, 2018

	Hospital A
Inpatient Mortality, AUROC[1](95% CI)	
Deep learning 24 hours after admission	**0.95**(0.94-0.96)
Full feature enhanced baseline at 24 hours after admission	0.93(0.92-0.95)

← Deep learning
← Logistic Regression

Classical ML Remains Competitive in Some Domains

Google Research, Scalable & accurate deep learning with electronic health records, NPJ Digital Medicine, 2018

Inpatient Mortality, AUROC[1](95% CI)	Hospital A
Deep learning 24 hours after admission	**0.95**(0.94-0.96)
Full feature enhanced baseline at 24 hours after admission	0.93 (0.92-0.95)

← Deep learning
← Logistic Regression

Original Article

Deep learning of aftershock patterns following large earthquakes

Phoebe M. R. DeVries, Fernanda Viégas, Martin Wattenberg & Brendan J. Meade

Nature **560**, 632–634(2018) | Cite this article

19k Accesses | **48** Citations | **1003** Altmetric | Metrics

Deep Neural Network

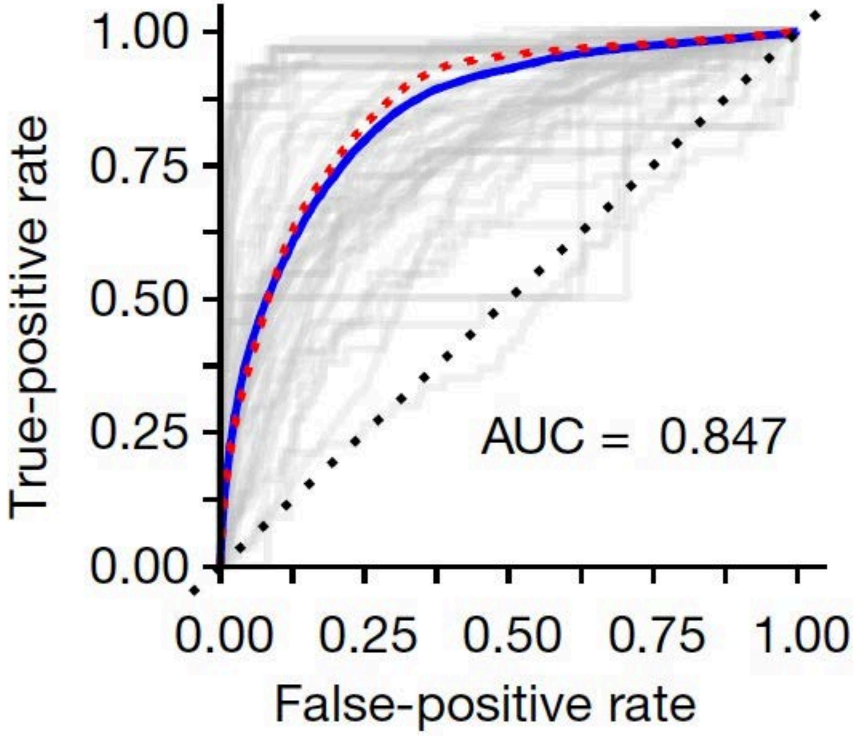

Classical ML Remains Competitive in Some Domains

Google Research, Scalable & accurate deep learning with electronic health records, NPJ Digital Medicine, 2018

Inpatient Mortality, AUROC[1](95% CI)	Hospital A	
Deep learning 24 hours after admission	**0.95**(0.94-0.96)	← Deep learning
Full feature enhanced baseline at 24 hours after admission	0.93 (0.92-0.95)	← Logistic Regression

Commentary

One neuron versus deep learning in aftershock prediction

Arnaud Mignan & Marco Broccardo

Nature **574**, E1–E3(2019) | Cite this article

7666 Accesses | **6** Citations | **356** Altmetric | Metrics

Logistic regression (2 input features)

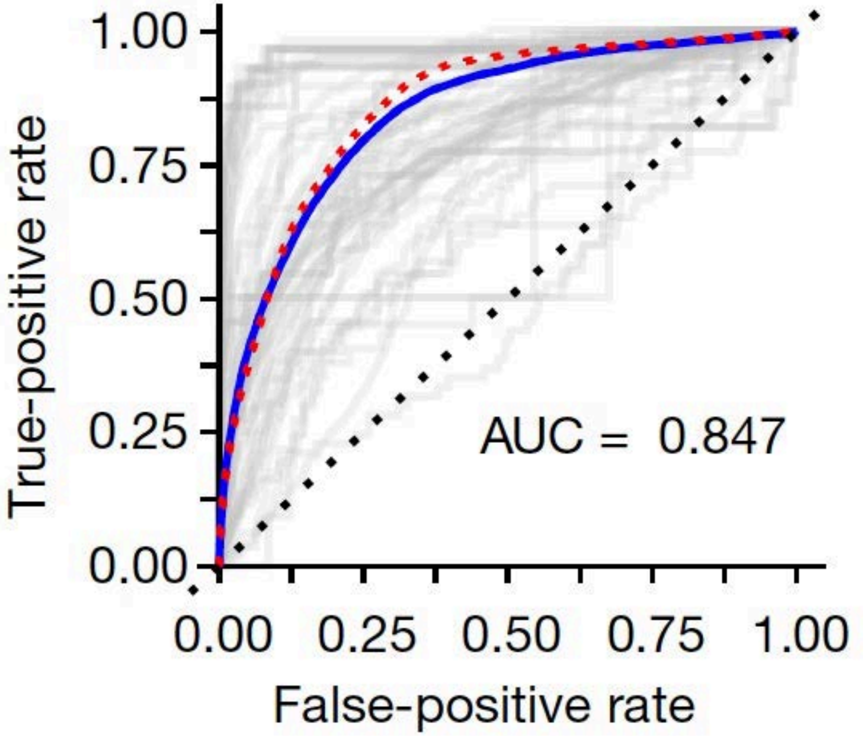

Applications of Deep Learning in Neuroscience

- Models of the brain
  - Vision (Khaligh-Razavi & Kriegeskorte, 2014; Yamins et al., 2014; Eickenberg et al., 2017; Bashivan et al., 2019)
  - Auditory perception (Kell et al., 2018)
  - Reinforcement learning (Dabney et al., 2020)
  - Grid cells / navigation (Banino et al., 2018)
- Tools to analyze brain data
  - Lesion/tumor segmentation (Pinto et al., 2016; Havaei et al., 2017; Kamnitsas et al., 2017b; Zhao et al., 2018)
  - Anatomical segmentation (Wachinger et al., 2018; Zhao et al., 2018)
  - Image modality/quality transfer (Bahrami et al., 2016; Nie et al., 2017; Blumberg
  - Image registration (Yang et al., 2017; Dalca et al., 2018)
  - Behavioral and disease prediction (Plis et al., 2014; van der Burgh et al., 2017; Vieira et al., 2017; Nguyen et al., 2018)

Evidence From Challenges/Competitions

- Advantage: test set is truly hidden

	TEAM	MAE	P
1	BrainAgeDifference *Donders Institute, Radboud University*	2.9043	-0.3914
2	BrainAGE *University Hospital Jena*	3.0857	-0.3423
3	ARAMIS *Brain and Spine Institute Paris*	3.3284	-0.2103
4	Quantum Pika *National Yang Ming University*	3.3315	-0.3939
5	sablab *Cornell University*	3.3716	-0.2469

Accurate brain age prediction with lightweight deep neural networks

Han Peng[1,2,3*†], Weikang Gong[1*], Christian F. Beckmann[1,3], Andrea Vedaldi[2], Stephen M. Smith[1]

[1]Centre for Functional MRI of the Brain (FMRIB), University of Oxford, Oxford, OX3 9DU, United Kingdom
[2]Visual Geometry Group (VGG), University of Oxford, Oxford, OX2 6NN, United Kingdom
[3]Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, 6525 EN, The Netherlands

PAC challenge 2019 (https://www.photon-ai.com/pac2019)

Evidence From Challenges/Competitions

- Advantage: test set is truly hidden

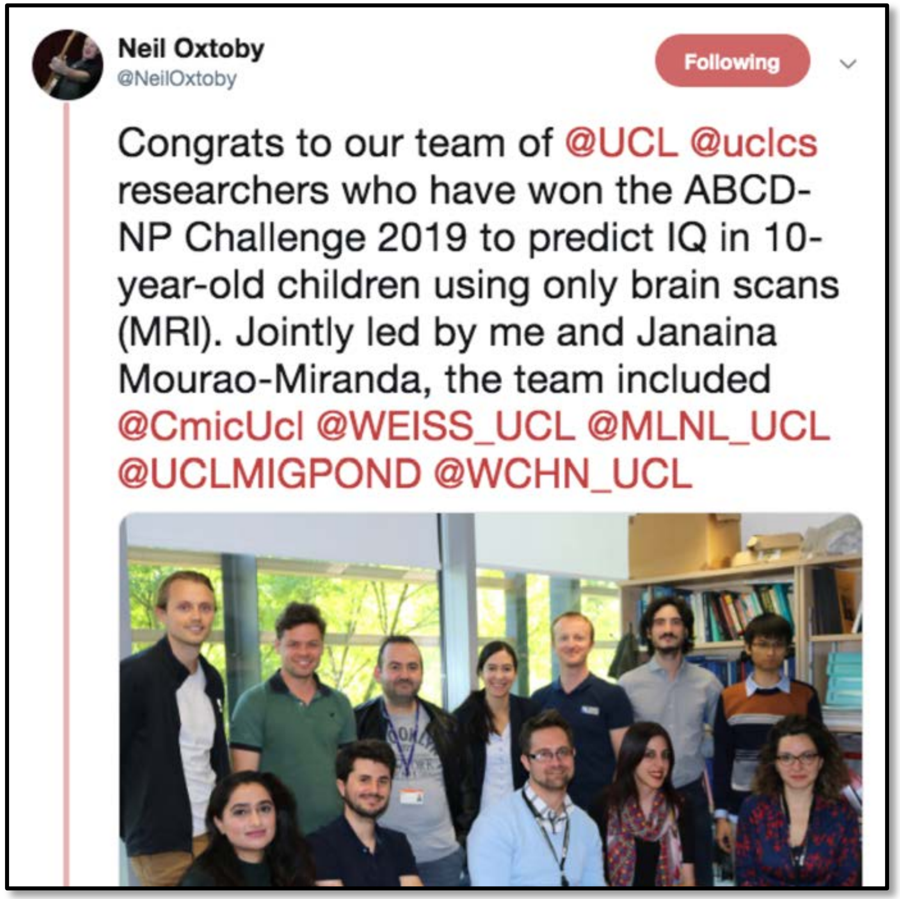

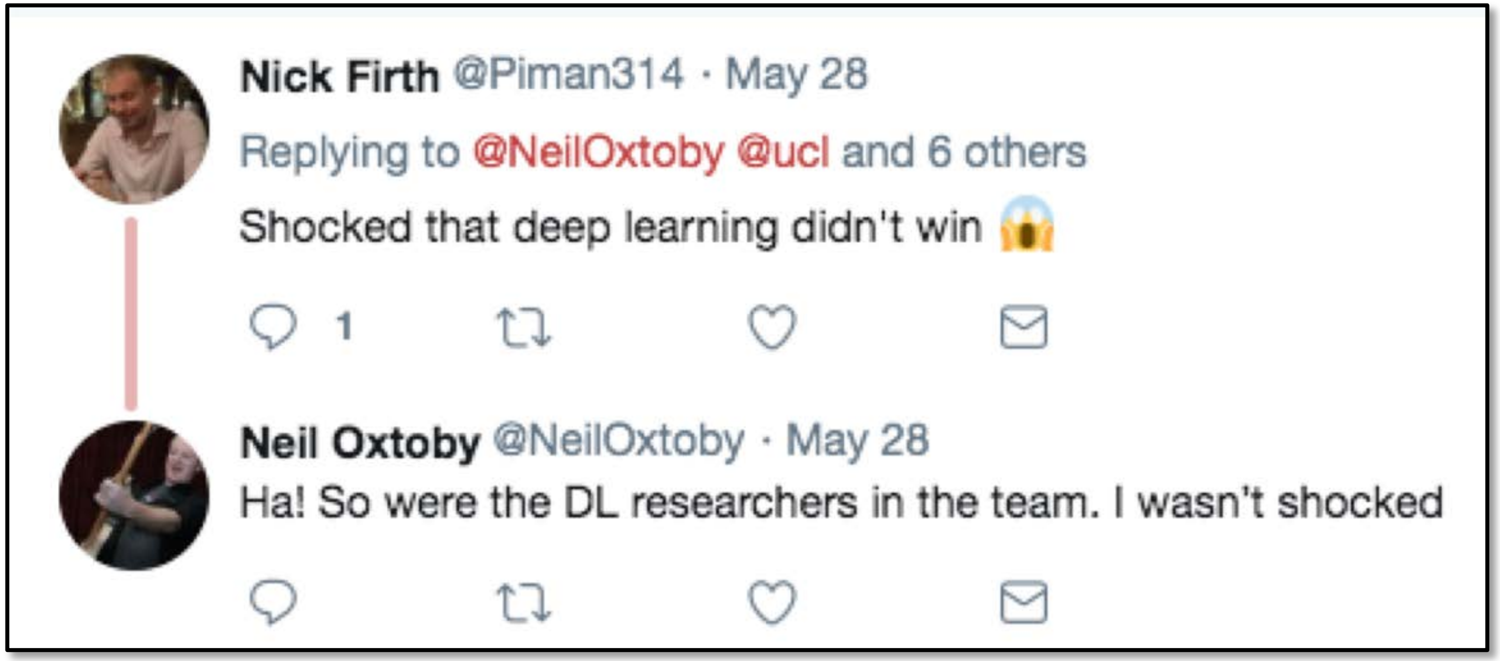

ABCD Fluid Intelligence Prediction Challenge 2019
(https://sibis.sri.com/abcd-np-challenge/)

Evidence From Challenges/Competitions

- Advantage: test set is truly hidden
- For anatomical/lesion segmentation challenges, deep learning has consistent winning track record (e.g., Choi et al., 2016; Kamnitsas et al., 2017, Li et al., 2018)
- For predicting behavior (e.g., fluid intelligence), age & diseases, deep learning has a more mixed record
    - PAC brain age prediction 2019: winning algorithm uses DNN
    - ABCD fluid intelligence prediction: 2019: winning algorithm uses kernel regression
    - TADPOLE challenge 2019: winning algorithm uses XGBoost
- Disadvantage: winner based on point estimate of prediction accuracy, so winner might not be statistically better than next best team => benchmarking studies are important

Evidence From Benchmarking Studies


Tong He

- ~10K subjects (UK Biobank + Human Connectome Project)
- Predict behavior & demographics with resting-state functional connectivity
- Kernel regression (KRR), fully-connected feedforward neural network (FNN), BrainNetCNN (Kawahara et al., 2017) & graph convolutional neural networks (GCNN)

Fluid Intelligence Prediction in UK Biobank
(Test Set N = 1000)

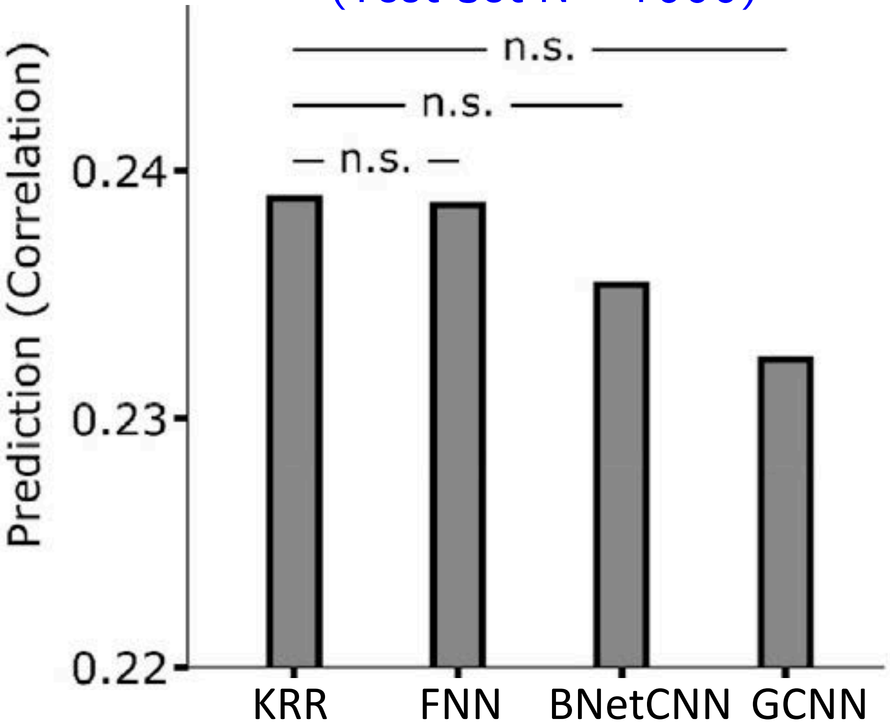

Sex Prediction in UK Biobank
(Test Set N = 1000)

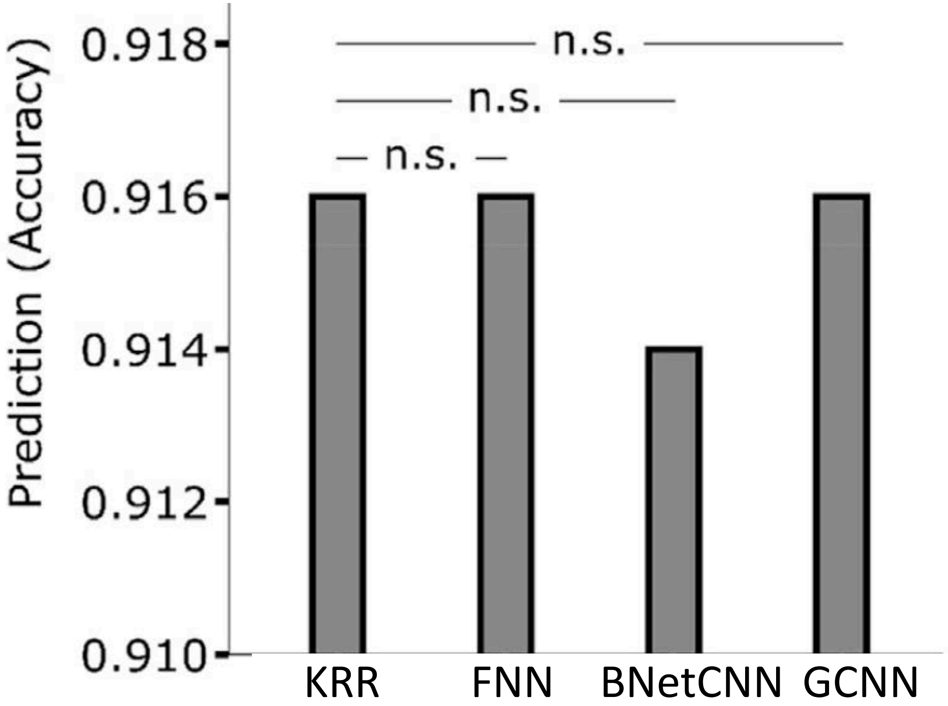

He, NeuroImage, 2020

Evidence From Benchmarking Studies

Tong He

- ~10K subjects (UK Biobank + Human Connectome Project)
- Predict behavior & demographics with resting-state functional connectivity
- Kernel regression (KRR), fully-connected feedforward neural network (FNN), BrainNetCNN (Kawahara et al., 2017) & graph convolutional neural networks (GCNN)

Fluid Intelligence Accuracy (Correlation) In Test Set (N = 1000)
Hyperparameters Tuned With Validation Set (N = 1000)

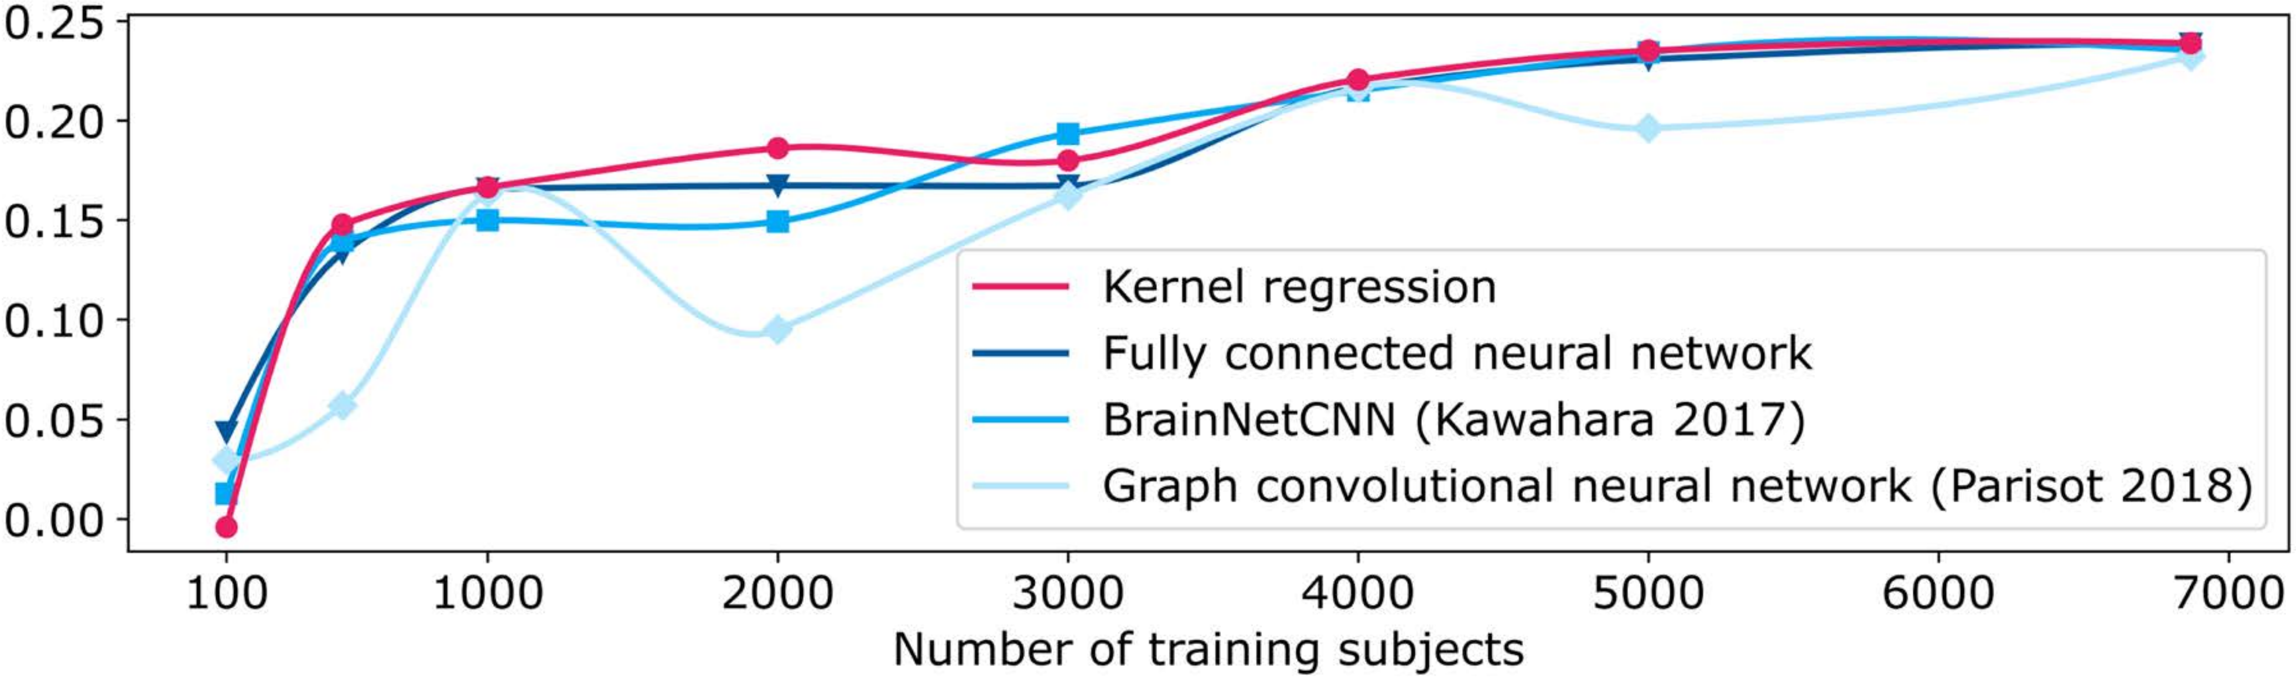

He, NeuroImage, 2020

Evidence From Benchmarking Studies

- ~14K subjects (UK Biobank + Human Connectome Project + ABIDE + ACPI)
- Predict using resting-state functional connectivity

additional shrinkage is not necessarily required on well regularized connectivity estimates in tangent space; however, could be applied for noiser data. Lastly, we evaluated various classifiers for prediction of non-imaging variables from connectivity estimates and **concluded that a carefully designed deep learning based architecture (2D RCNN) can be a valuable tool for analyzing functional connectivity. However, Elastic Net probably performs better at present overall.**

Pervaiz et al., NeuroImage, 2020

Evidence From Benchmarking Studies

- Hypothesis: hard to exploit nonlinearity in MRI data?
  - Computer Vision: Classical Linear < Classical Nonlinear < Deep Neural Networks

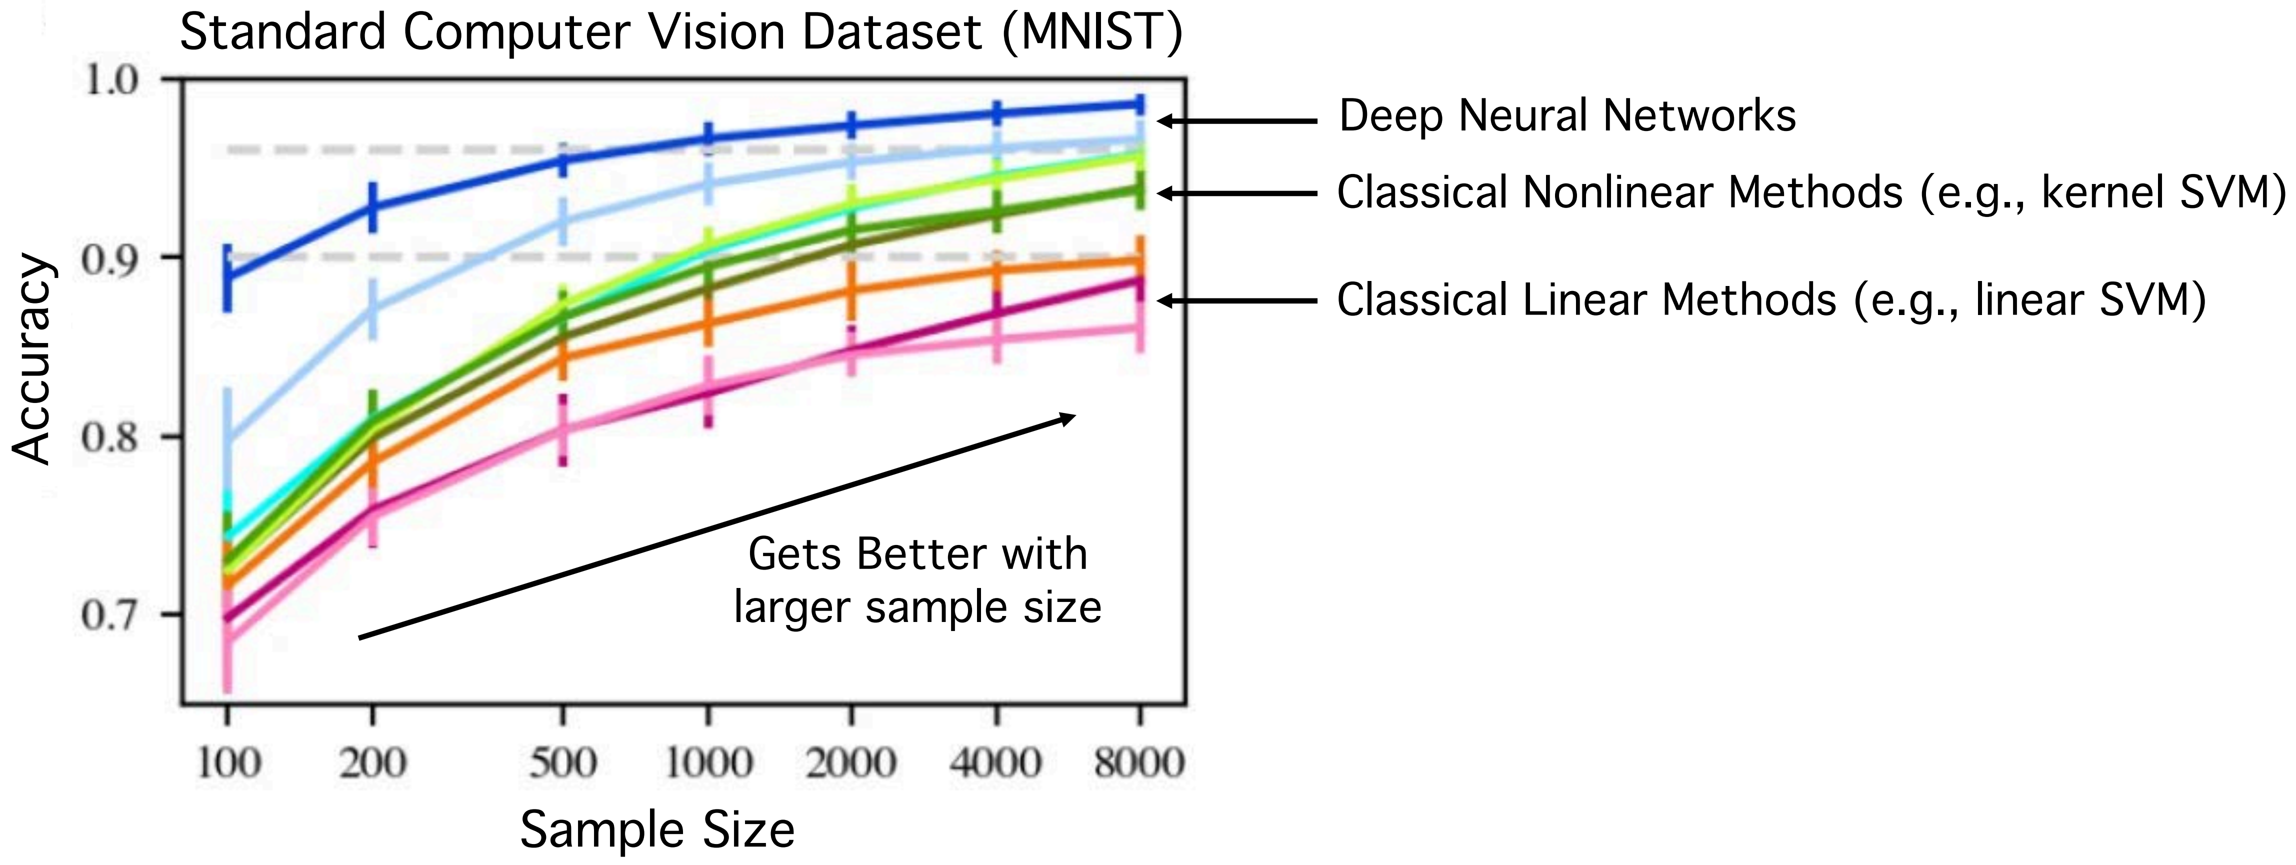

Schulz et al., biorxiv

Evidence From Benchmarking Studies

- Hypothesis: hard to exploit nonlinearity in MRI data?
  - Computer Vision: Classical Linear < Classical Nonlinear < Deep Neural Networks
  - MRI: Almost equivalent performance across all methods

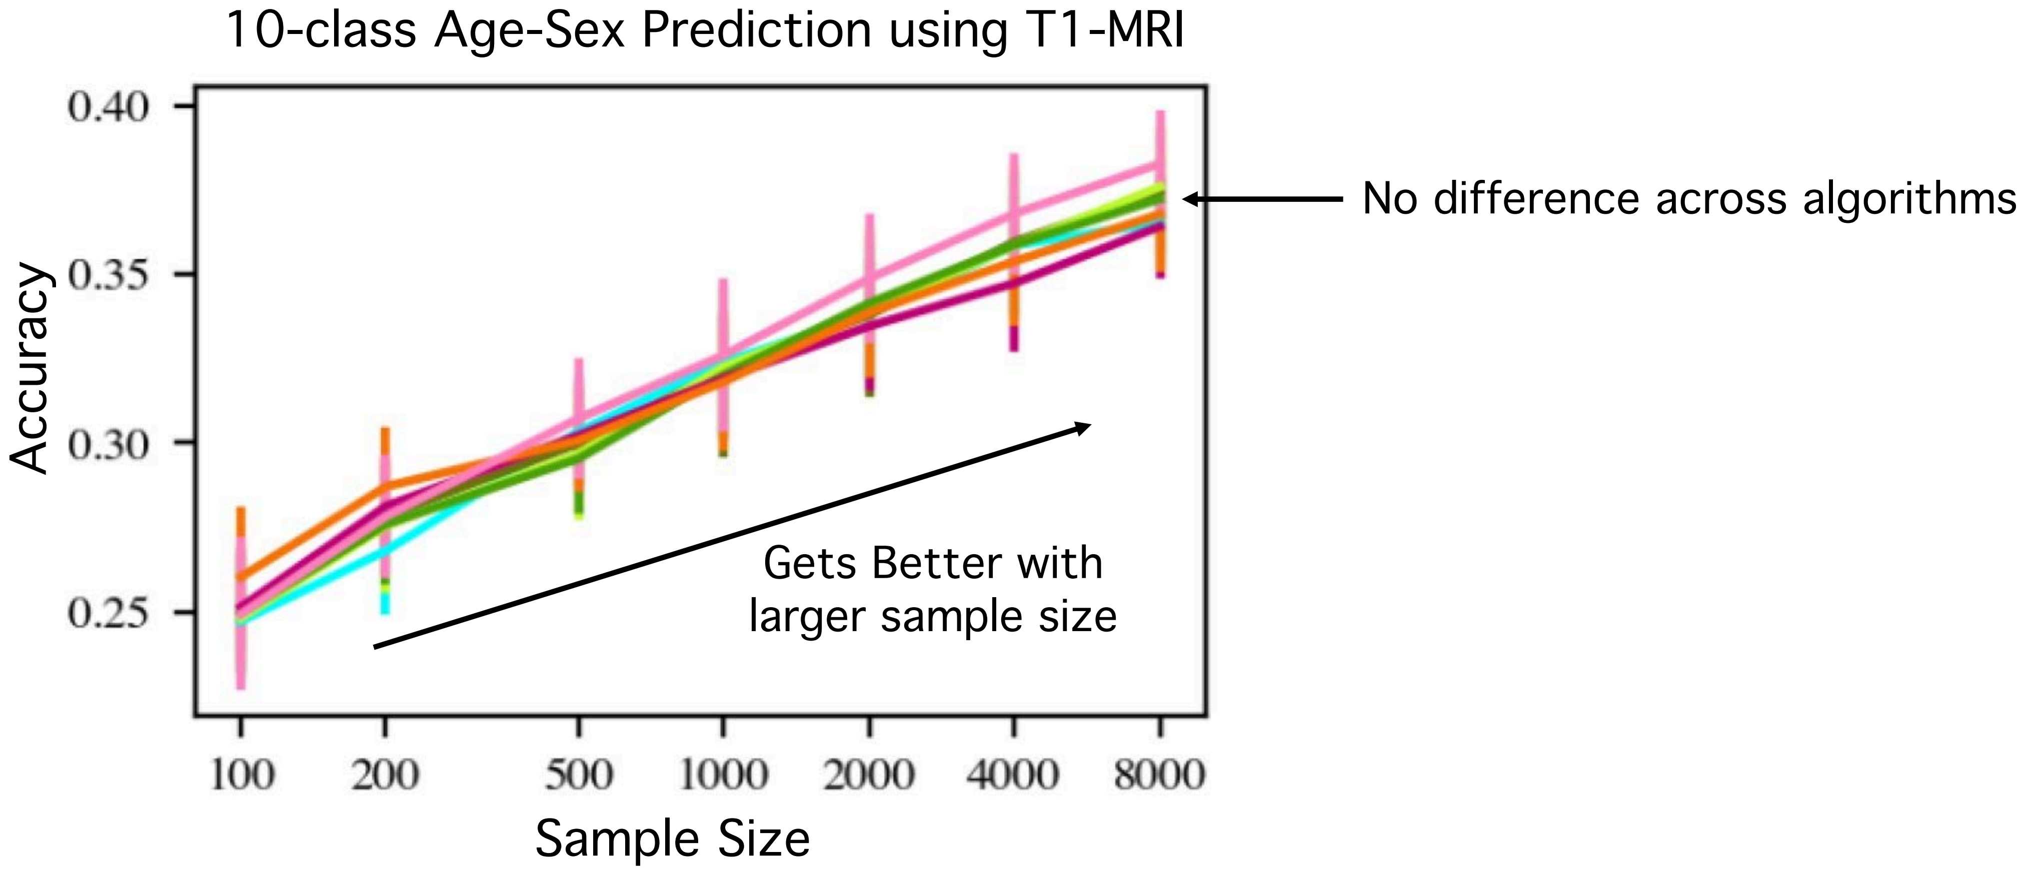

Schulz et al., biorxiv

Cautionary Tale

Minh Nguyen

- Not claiming deep learning is not useful (we also use DNNs)
- TADPOLE Challenge (ADNI)
  - Given 23 multimodal biomarkers at one or more timepoints
  - Predict cognitive scores, brain atrophy, clinical diagnosis for every month (indefinitely) into the future
  - Evaluation done for available time points (~5 years in the dataset)
- Challenge: lots of missing data
  - CDR available in 70% timepoints
  - Flortaucipir PET available in 16% timepoints
- Most studies assume all data present
  - Lose a lot of data
  - Not realistic in practice
- Our strategy: deep recurrent neural network (RNN) can handle missing data and model individual's longitudinal data

RNN Outperforms Baseline Algorithms

- 10-fold Nested Cross-Validation
- Evaluate prediction of diagnosis, ventricular volume, cognition (ADAS)

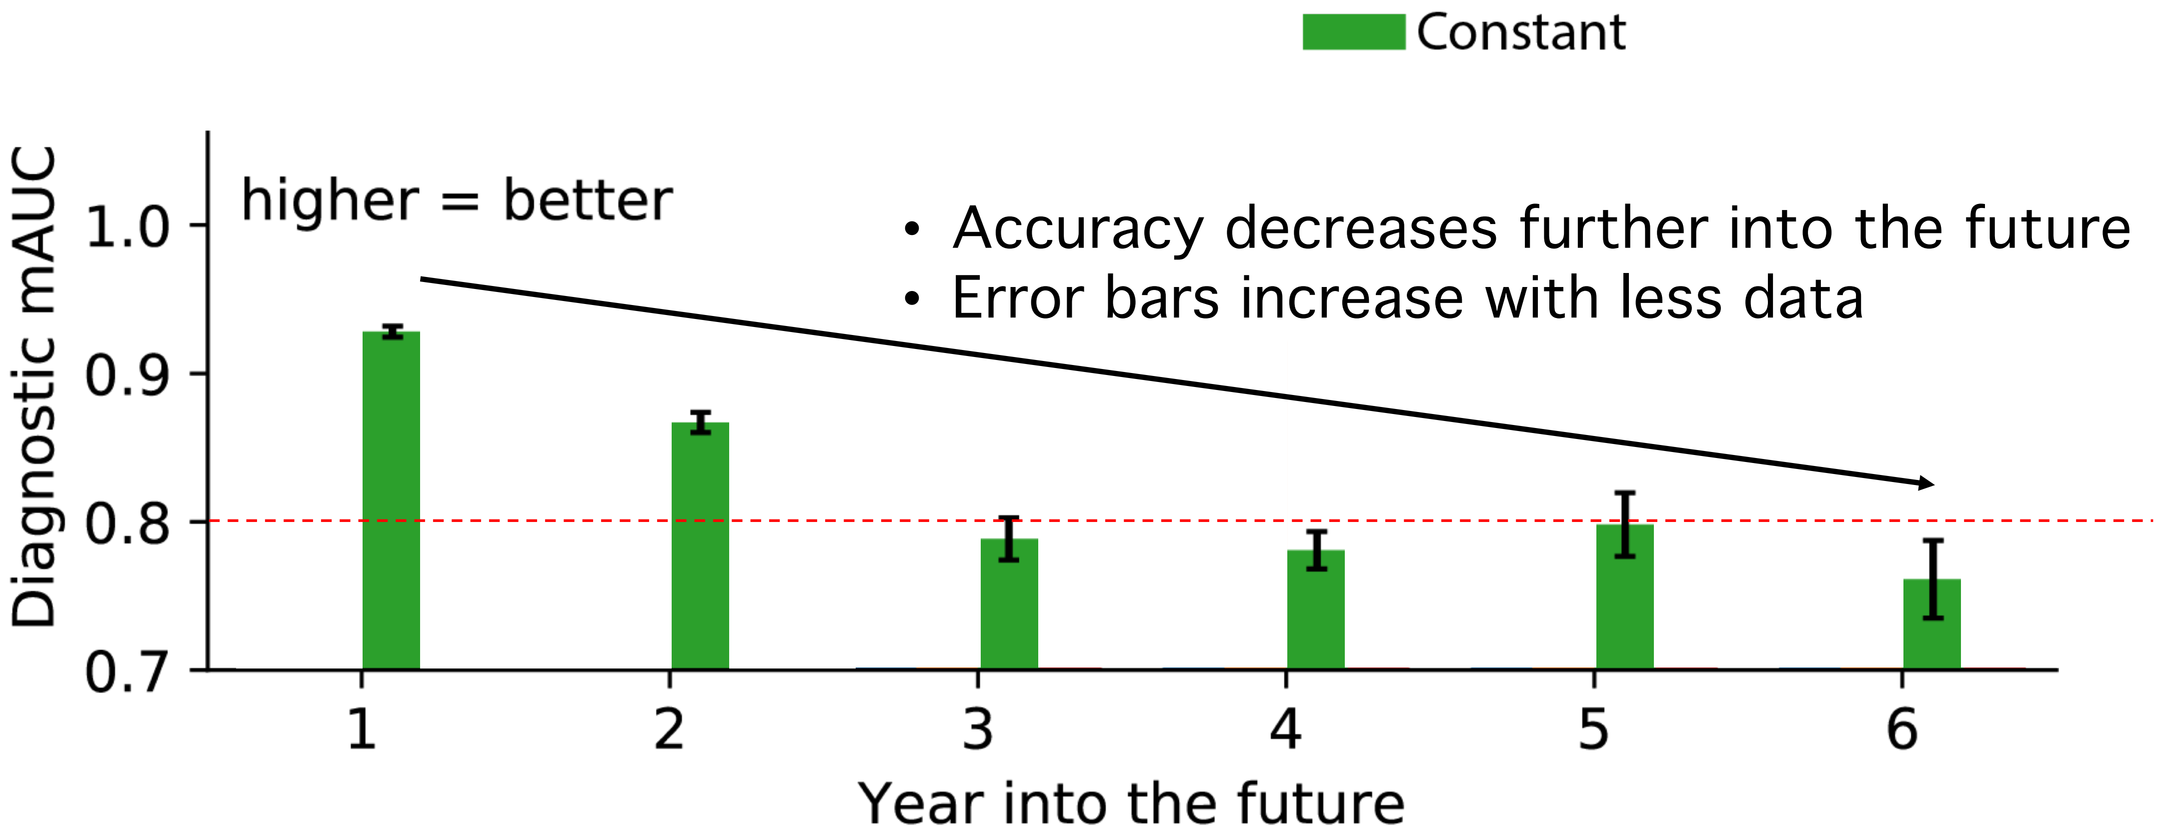

- Accuracy decreases further into the future
- Error bars increase with less data

Nguyen et al., PRNI, 2018
Nguyen et al., under revision

RNN Outperforms Baseline Algorithms

- 10-fold Nested Cross-Validation
- Evaluate prediction of diagnosis, ventricular volume, cognition (ADAS)

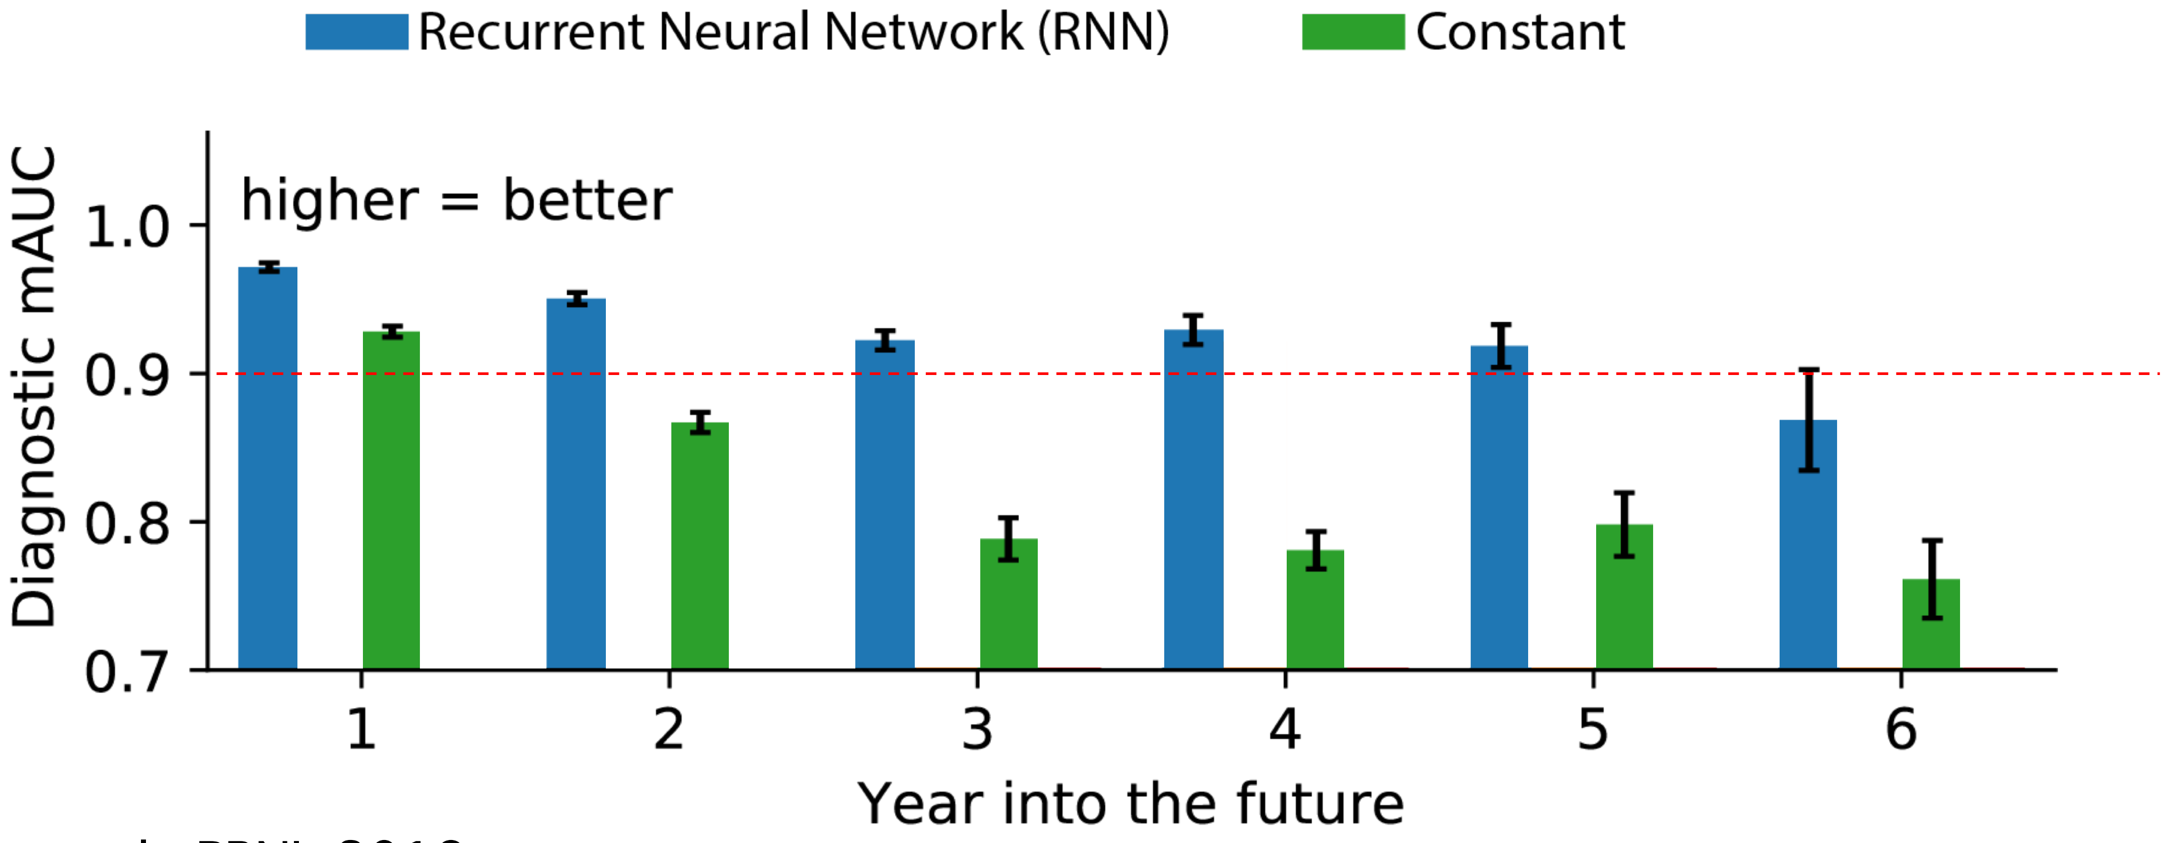

Nguyen et al., PRNI, 2018
Nguyen et al., under revision

RNN Outperforms Baseline Algorithms

- 10-fold Nested Cross-Validation
- Evaluate prediction of diagnosis, ventricular volume, cognition (ADAS)

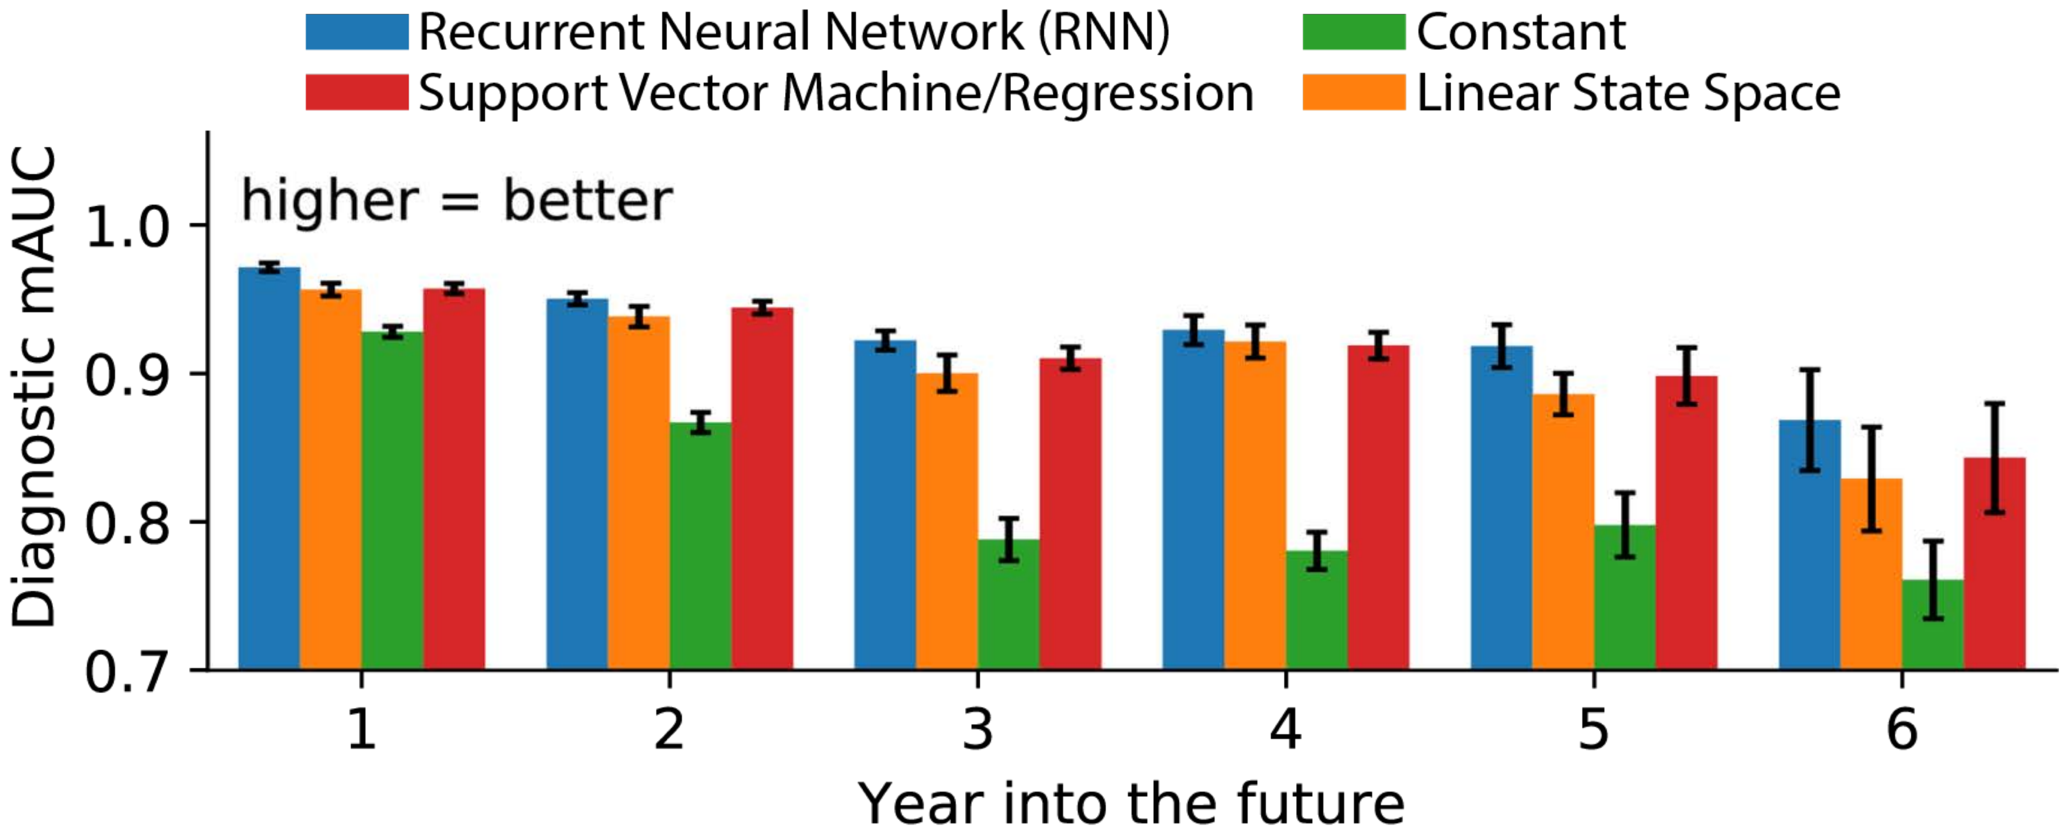

Nguyen et al., PRNI, 2018
Nguyen et al., under revision

TADPOLE Leaderboard

- Completely unseen data from ADNI
- 5th in the original challenge
- Currently 2nd out of 63 entries (as of June 3rd, 2020)

RANK	FILE NAME	MAUC RANK	MAUC	BCA	ADAS RANK	ADAS MAE	ADAS WES	ADAS CPA	VENTS RANK	VENTS MAE	VENTS WES	VENTS CPA
1.0	Frog	1.0	0.931	0.849	5.0	4.85	4.74	0.44	10.0	0.45	0.33	0.47
2.0	CBIL-MinMFa	3.0	0.909	0.845	16.0	5.53	5.54	0.39	13.0	0.46	0.46	0.02
3.0	CBIL-MinMF1	15.0	0.886	0.818	7.0	5.10	5.11	0.41	14.0	0.46	0.46	0.37
4.0	EMC1-Std	9.0	0.898	0.811	27.5	6.05	5.40	0.45	1.5	0.41	0.29	0.43
5.0	VikingAI-Sigmoid	21.0	0.875	0.760	9.0	5.20	5.11	0.02	11.5	0.45	0.35	0.20
6.0	EMC1-Custom	13.0	0.892	0.798	27.5	6.05	5.40	0.45	1.5	0.41	0.29	0.43

XGBoost

https://tadpole.grand-challenge.org/D4_Leaderboard/

Tips When Reading DNN Papers

- Is it hype?
  - If “deep learning” replaced with “logistic regression” in paper, is it still exciting?
  - Why is it advantageous to use this particular DNN?
- Sample size
  - Even 100 subjects => ±10% error bars

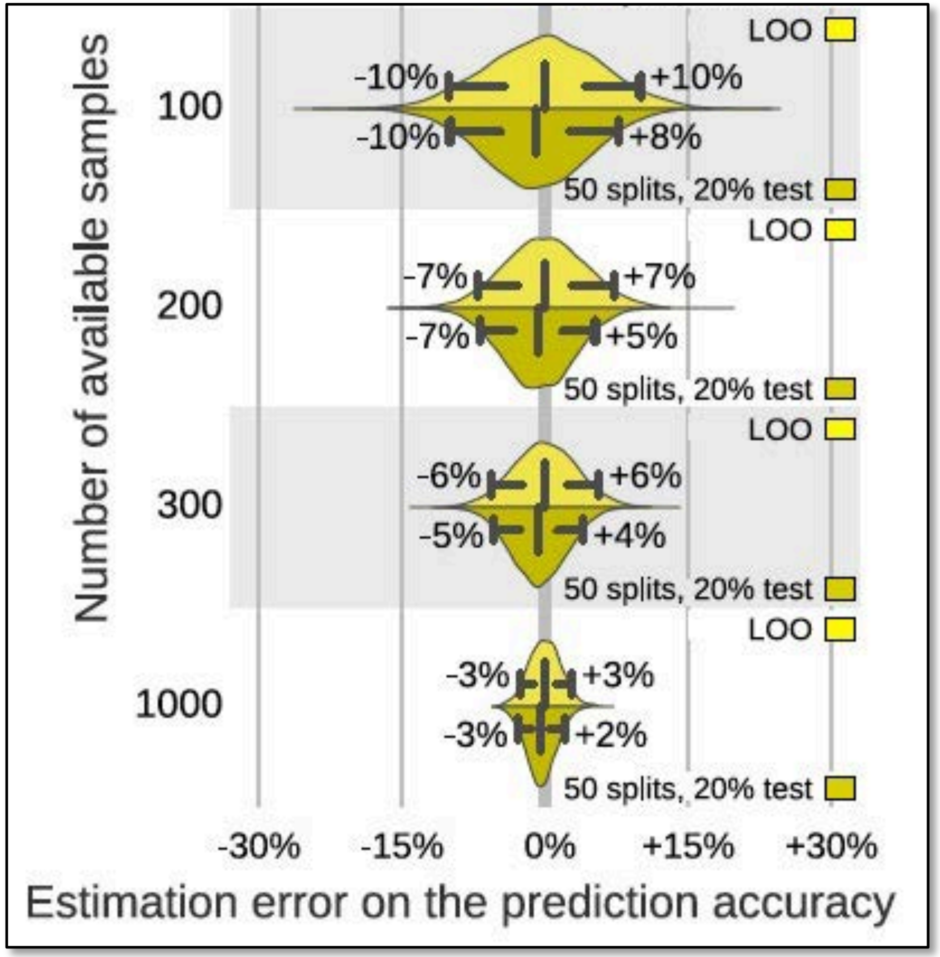

Varoquaux, NeuroImage, 2018

Tips When Reading DNN Papers

- Is it hype?
  - If "deep learning" replaced with "logistic regression" in paper, is it still exciting?
  - Why is it advantageous to use this particular DNN?
- Sample size
  - Even 100 subjects => ±10% error bars
  - >>100 subjects for DNNs to be credible
- Many hyperparameters in DNNs
  - Is it clear how the authors optimize their hyperparameters?
  - If authors manually tune hyperparameters, then inner-loop (nested) cross-validation is bogus because information from tuning one fold will leak to another fold (via the person tuning the hyperparameters)
  - Recommended scheme: training, validation, test
    - Training set to train model
    - Validation set to tune hyperparameters
    - Test set to perform final evaluation

Summary

- DNNs have revolutionized machine learning
- As models of the brain, DNNs have provided new mechanistic insights
- As tools for analyzing brain data, DNNs is promising
- DNNs excel in anatomical/lesion/tumor segmentation, image registration, image modality/quality transfer
  - Perhaps convolutional neural networks can better exploit nonlinearity in 3D ”structural” data for “image-processing-type” problems?
- Currently still unclear about DNNs’ advantage when predicting behavior/age/disease with anatomical T1 or fMRI
  - Mixed results from competitions/challenges
  - 3 Benchmarking studies (He et al., NeuroImage, 2020; Pervaiz et al., NeuroImage, 2020; Schulz et al., biorxiv, 2020) suggest comparable performance between DNNs and classical approaches
  - 1 Benchmarking study (Abrol et al., biorxiv, 2020) suggest that other benchmarking studies not making good use of DNNs

Funding & Support

Ministry of Education
SINGAPORE

NATIONAL RESEARCH FOUNDATION
PRIME MINISTER'S OFFICE
SINGAPORE

SG NSCC SINGAPORE
National Supercomputing Centre

CENTRE for SLEEP & COGNITION

NUS
National University of Singapore

NMRC National Medical Research Council
Singapore

CLINICAL IMAGING RESEARCH CENTRE
SINGAPORE
A joint venture between the Agency for Science, Technology And Research (A*STAR) and the National University of Singapore (NUS)

NUS National University of Singapore | N.1 INSTITUTE FOR HEALTH

Lab Members

Jianzhong Chen, Lijun An, Tong He, Ruby Kong, Xiaolu Kong, Valeria Kebets, Jingwei Li, Minh Binh Nguyen, Leon Ooi, Csaba Orban, Angela Tam, Aihuiping Xue, Xiaoxuan Yan, Shaoshi Zhang

Collaborators

Kevin Anderson
Boris Bernhardt
Maxwell Bertolero
Randy Buckner
Danilo Bzdok
Michael Chee
Gustavo Deco
Adriana Di Martino
Mark D'Esposito
Simon Eickhoff
Dorothea Floris
Peter Fox
Tian Ge
Sarah Genon
Evan Gordon
Avram Holmes

Timothy Laumann
Hesheng Liu
Beth Mormino
John Murray
Rik Ossenkoppele
Russ Poldrack
Mert Sabuncu
Nathan Spreng
Avi Snyder
Lucina Uddin
Dimtri Van De Ville
Sofie Valk
M. van den Heuvel
Xi-Nian Zuo
Helen Zhou

Women in Neuroscience

Repository for Women in Neuroscience

- www.winrepo.org
- over 1,100 profiles
- easy search
- recommendations

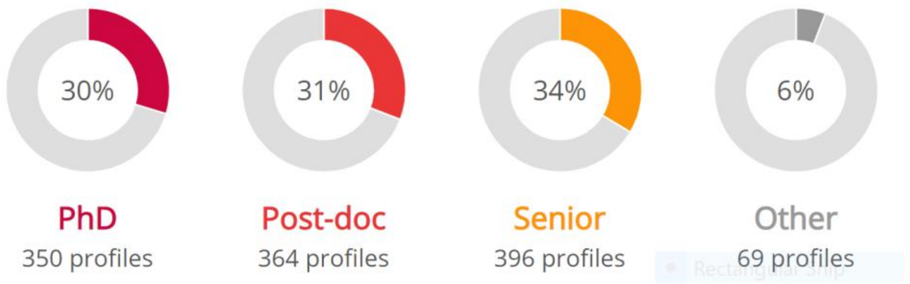

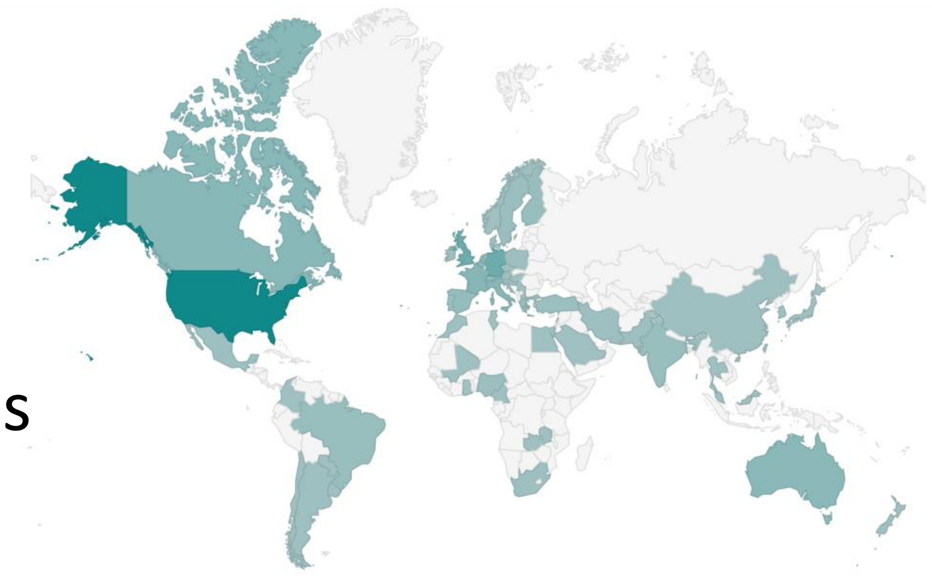

Support the project:

➢ sign up
➢ spread the word
➢ submit recommendations

WINRePo1
www.facebook.com/WiNRepository/