Learning from fMRI using Recurrent Neural Networks with Applications in Autism

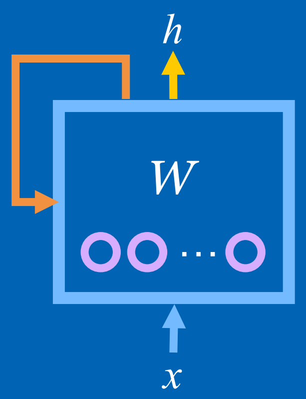

Nicha C. Dvornek
Associate Research Scientist, Radiology & Biomedical Imaging
MRRC fMRI Seminar Series
May 3, 2018

Questions to answer

- Why supervised machine learning for fMRI analysis?
- What is / why use deep learning with recurrent neural networks?
- How can we apply recurrent neural networks to fMRI analysis?

Outline

1. Supervised machine learning for fMRI
2. Deep learning and recurrent neural networks
3. Classification of autism / control from
   - Resting-state fMRI (rsfMRI)
   - rsfMRI + phenotypic data
4. Prediction of autism treatment outcome
5. Final thoughts

Traditional fMRI analysis finds *descriptive* model

- Use all the data to fit a model
- General linear model for task fMRI (tfMRI)

Each voxel:

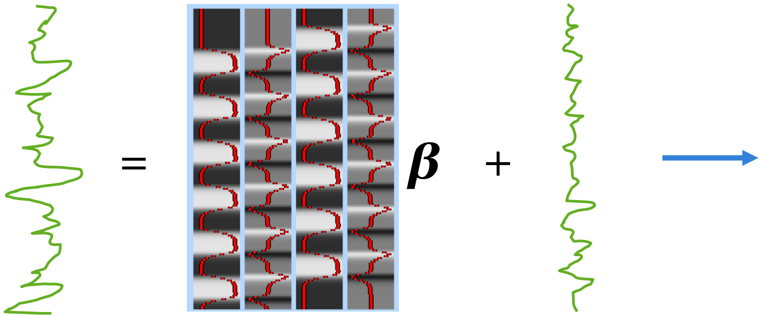

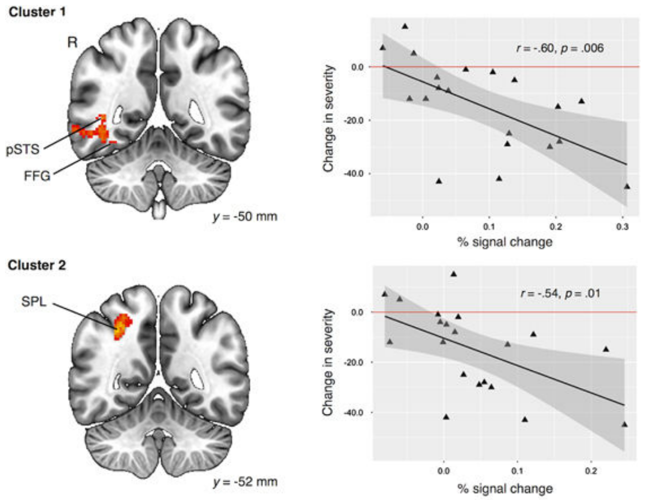

[1]

[1]Yang et al., Transl. Psychiatry 2017

Traditional fMRI analysis finds *descriptive* model

- Functional connectivity analysis for resting-state fMRI (rsfMRI)
- Use all the data to fit a model

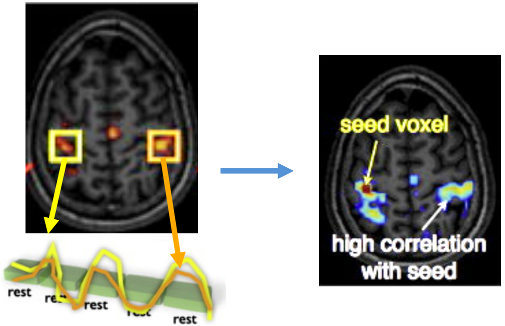

[1]

Seed-based correlations

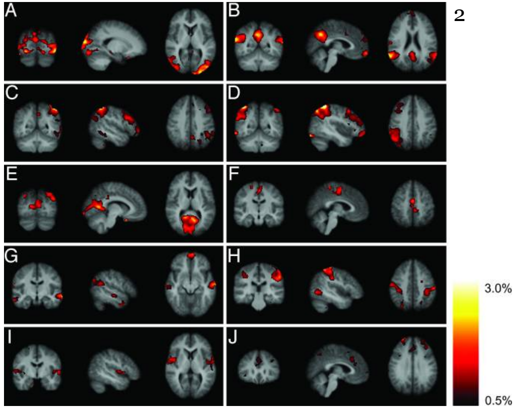

[2]

Independent component analysis

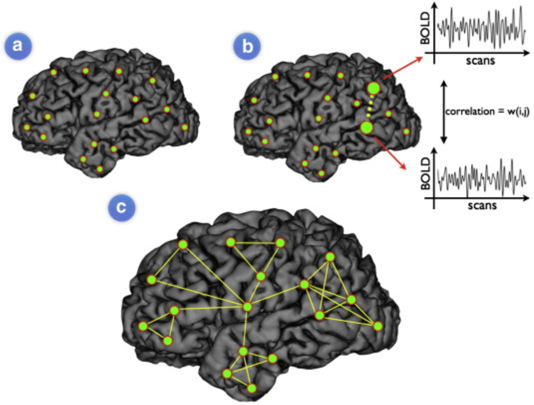

[1]

Graph-based analysis

[1]van der Heuvel et al., Eur. Neuropsychopharmacol. 2010. [2]Damoiseaux et al., PNAS 2006.

Supervised machine learning finds *predictive* model

- Training phase:

Inputs X → Learning Algorithm ? → Learned Predictive Model $f : X \rightarrow Y$

Outputs Y → Learning Algorithm

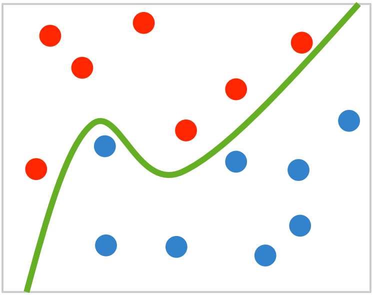

Supervised machine learning finds *predictive* model

- Testing phase:

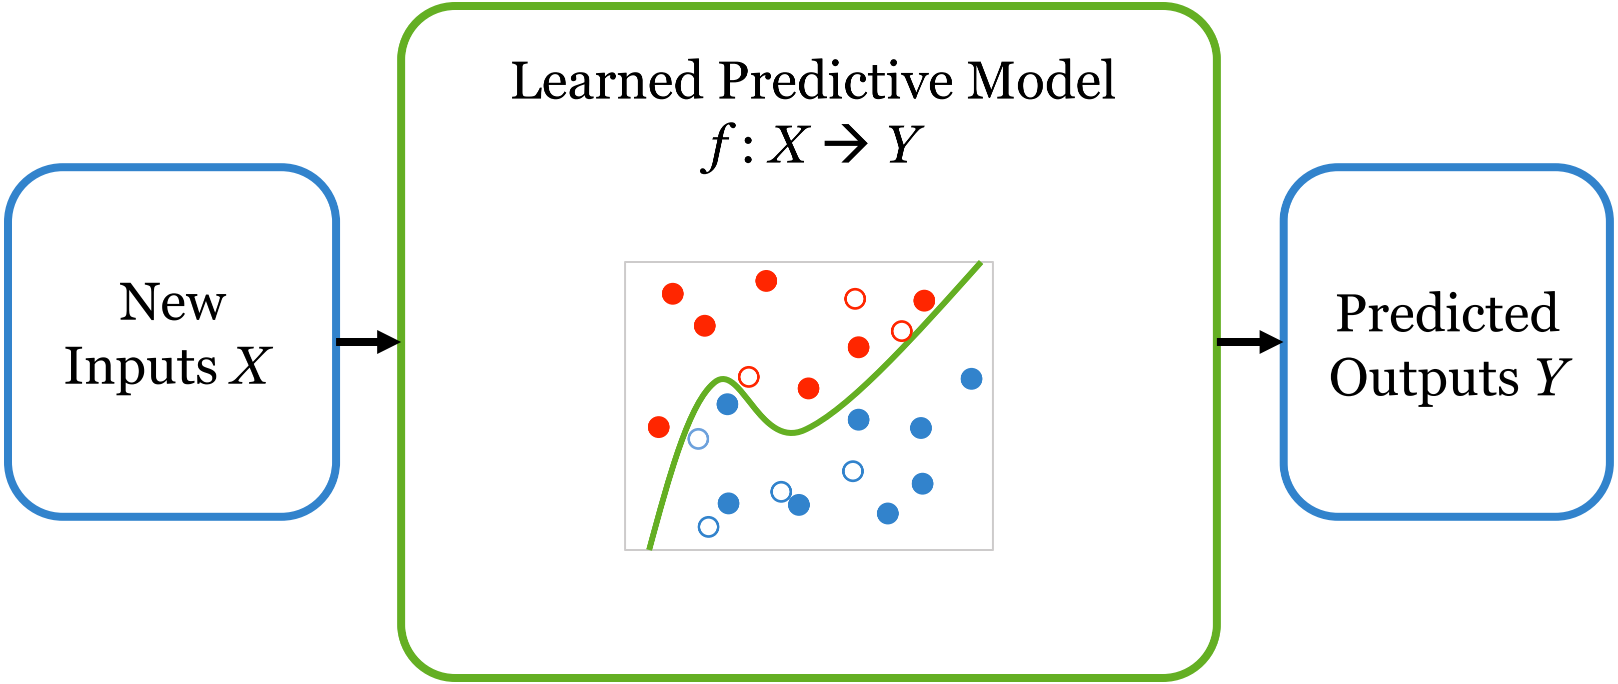

Application 1: Identify fMRI biomarkers for autism spectrum disorder (ASD)

1. Train classifier for ASD   2. Extract biomarkers from learned model

Inputs X

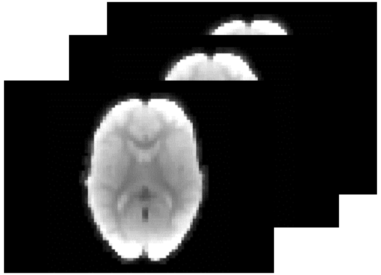

Outputs Y

ASD or TDC

→ **Learning Algorithm**

?

→ **Learned Predictive Model**

$f : X \rightarrow Y$

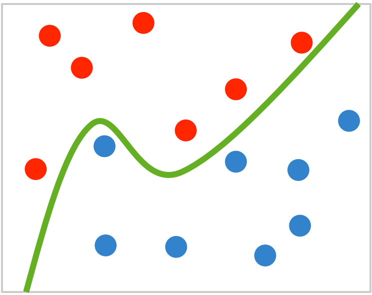

Application 2: Predict treatment outcome from baseline fMRI

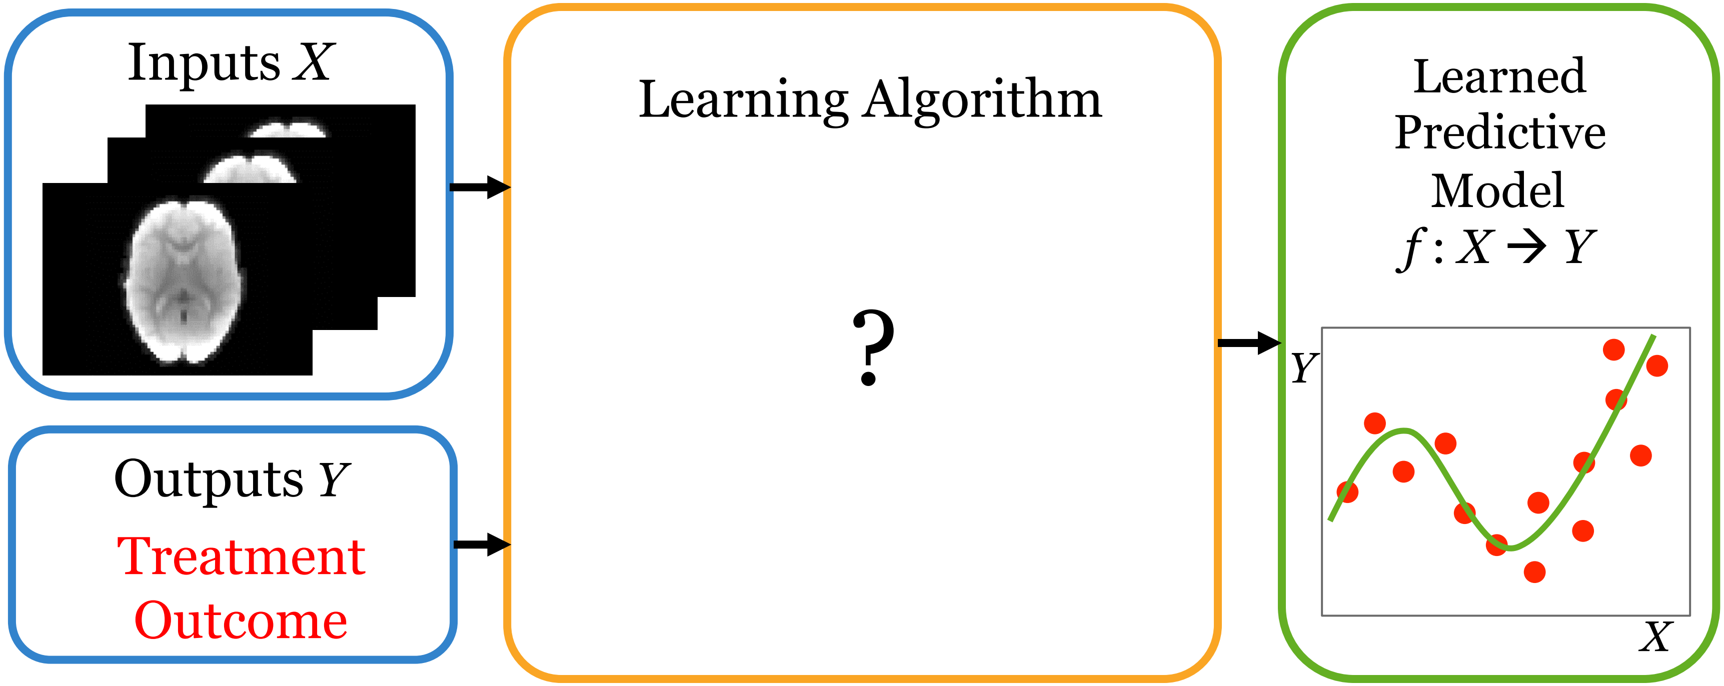

Outline

1. Supervised machine learning for fMRI
2. Deep learning and recurrent neural networks
3. Classification of autism / control from
    - Resting-state fMRI (rsfMRI)
    - rsfMRI + phenotypic data
4. Prediction of autism treatment outcome
5. Final thoughts

Deep learning uses artificial neural network model with many layers

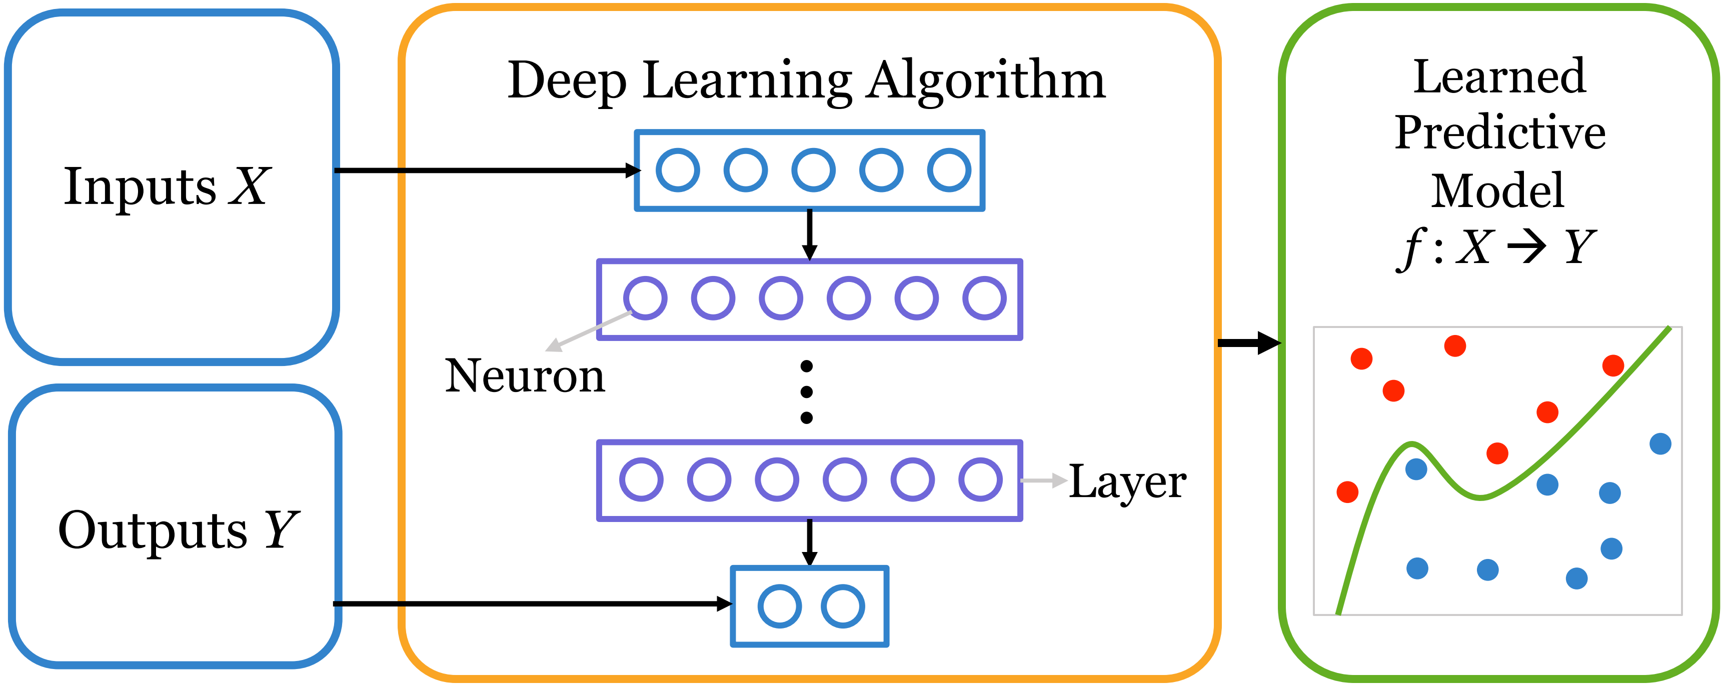

Why has deep learning become so successful/popular?

- Big computational resources
- Big datasets
- Big models

Why has deep learning become so successful/popular?

- Big computational resources
- Big datasets
- Big models
- It works!

Classification

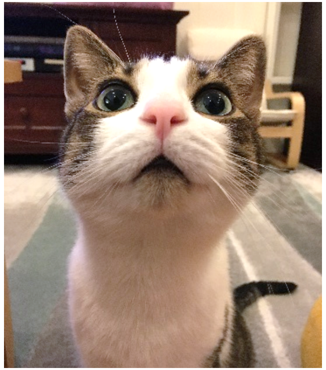

Cat

Detection

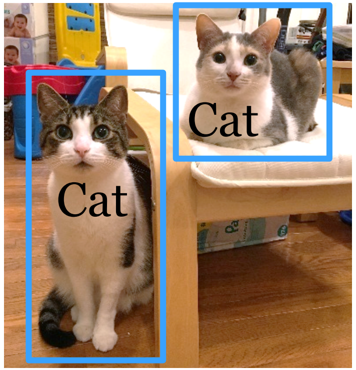

Captioning

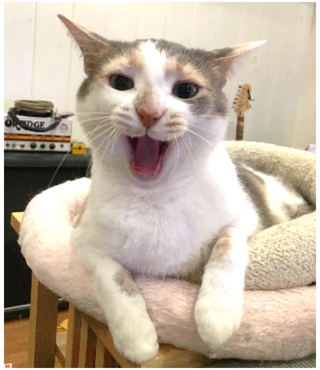

A cat is yawning

Recurrent neural networks were designed for temporal data

- Deep learning model with recurrent connections
  - "Memory" is passed forward

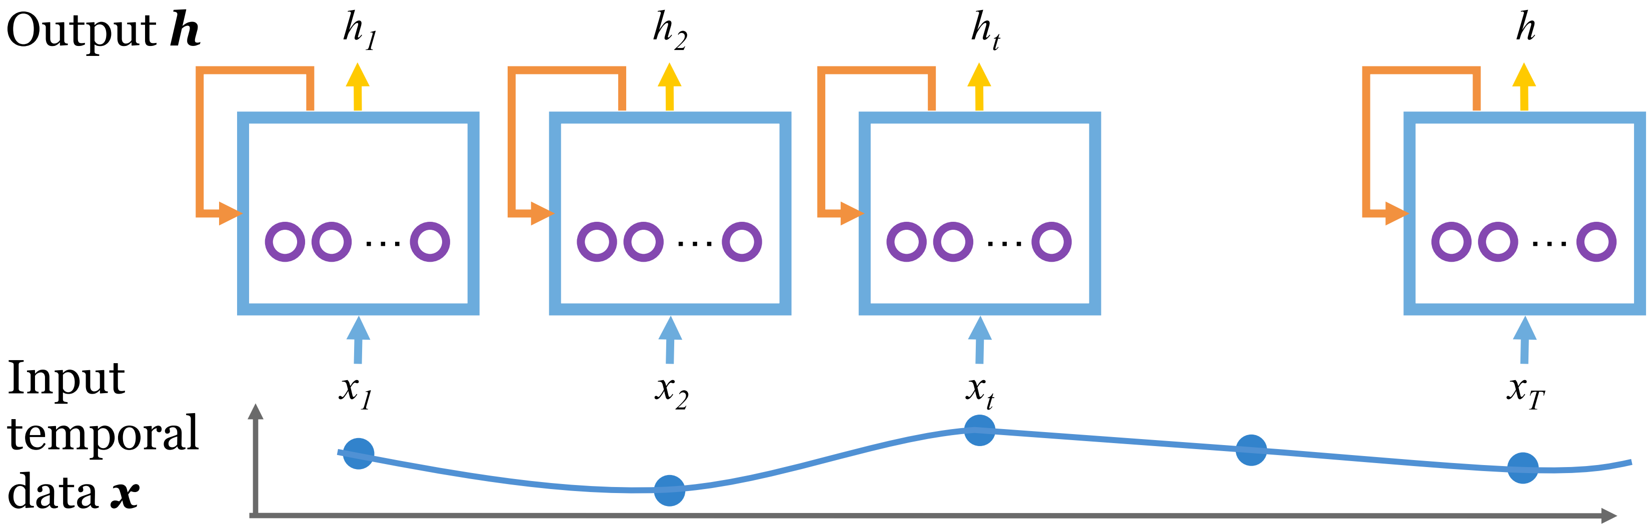

Recurrent neural networks were designed for temporal data

- Deep learning model with recurrent connections
  - "Memory" is passed forward
  - Model parameters W shared across time

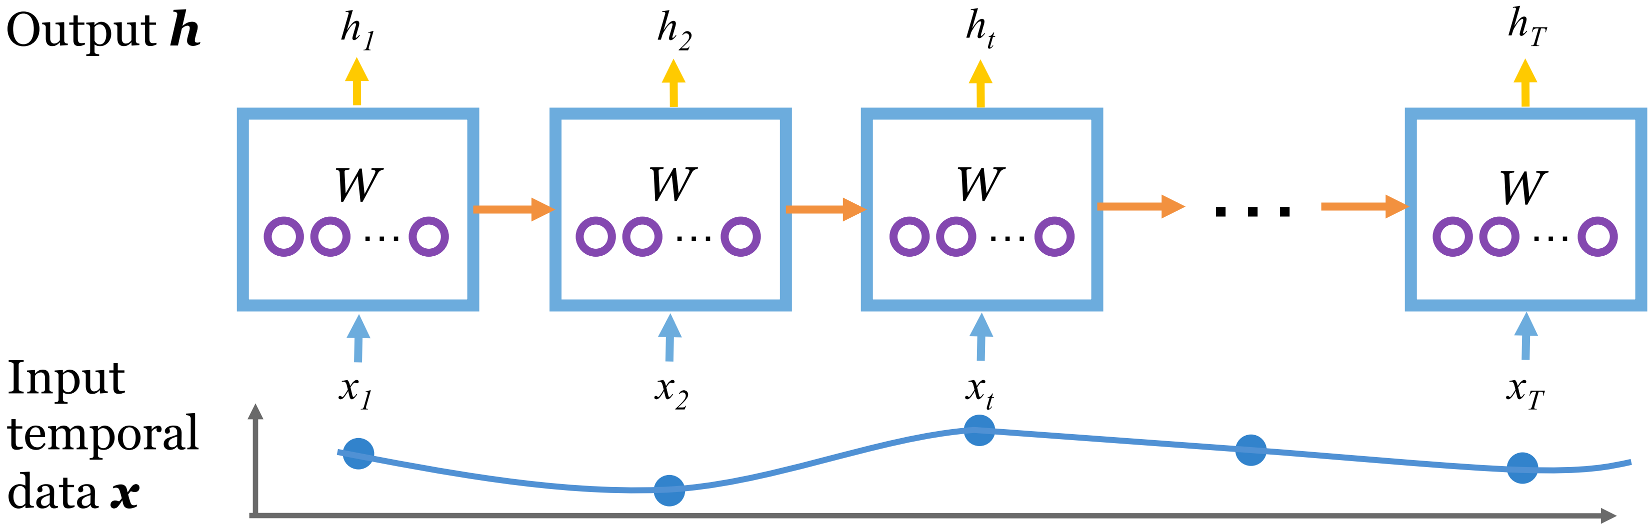

Recurrent neural networks with long short-term memory (LSTM) can learn long-term dependencies

- 2 types of recurrent information
    1. Hidden state h_t (output)
    2. Cell state c_t (remembers long-term information)

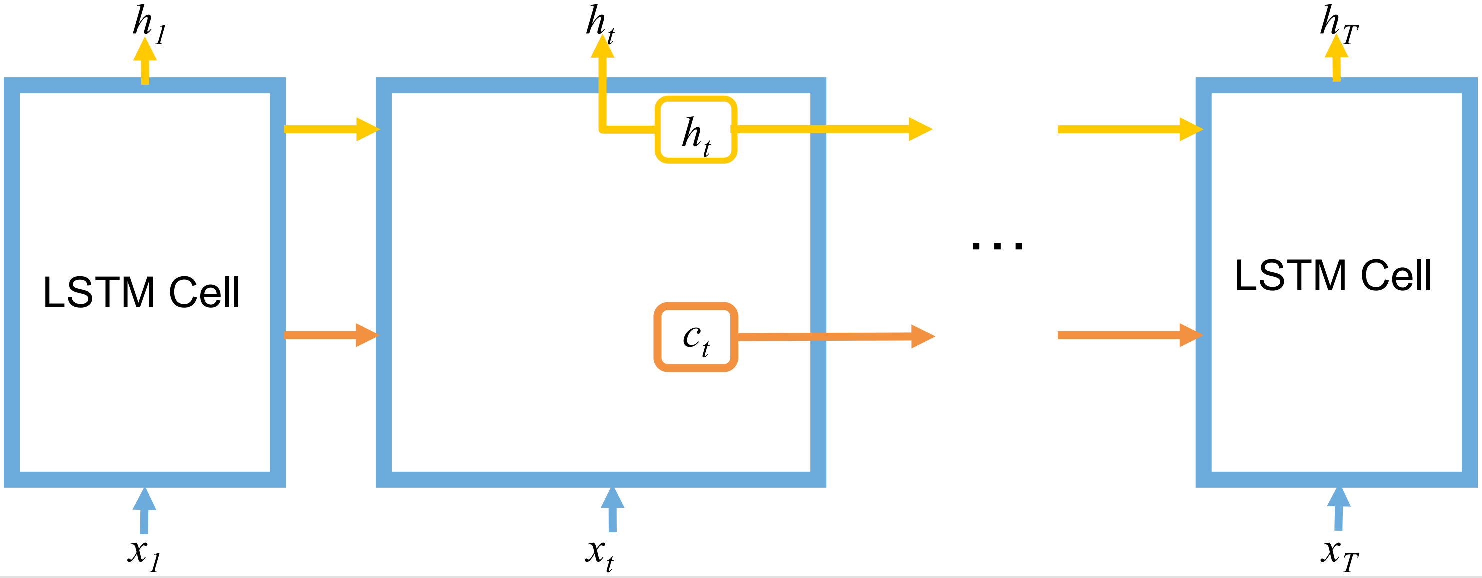

Recurrent neural networks with long short-term memory (LSTM) can learn long-term dependencies

- 4 neural network layers take x_t, h_{t-1} as inputs
  1. Estimated cell state $\tilde{c}_t$
  2. Input gate i_t
  3. Forget gate f_t
  4. Output gate o_t

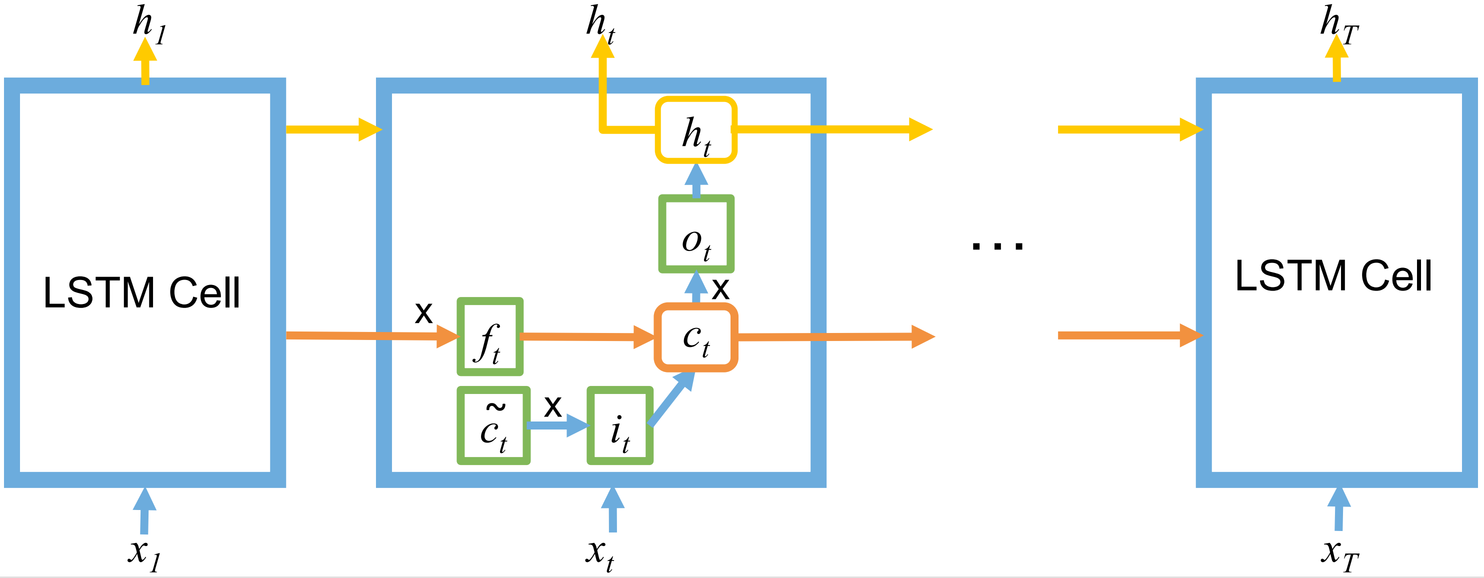

LSTM Equations

input gate $$i_t = \sigma\left(W_i x_t + U_i h_{t-1} + b_i\right)$$

forget gate $$f_t = \sigma\left(W_f x_t + U_f h_{t-1} + b_f\right)$$

current cell state $$\tilde{c}_t = \tanh\left(W_c x_t + U_c h_{t-1} + b_c\right)$$

cell state $$c_t = i_t * \tilde{c}_t + f_t * c_{t-1}$$

output gate $$o_t = \sigma\left(W_o x_t + U_o h_{t-1} + b_o\right)$$

hidden state $$h_t = o_t * \tanh\left(c_t\right)$$

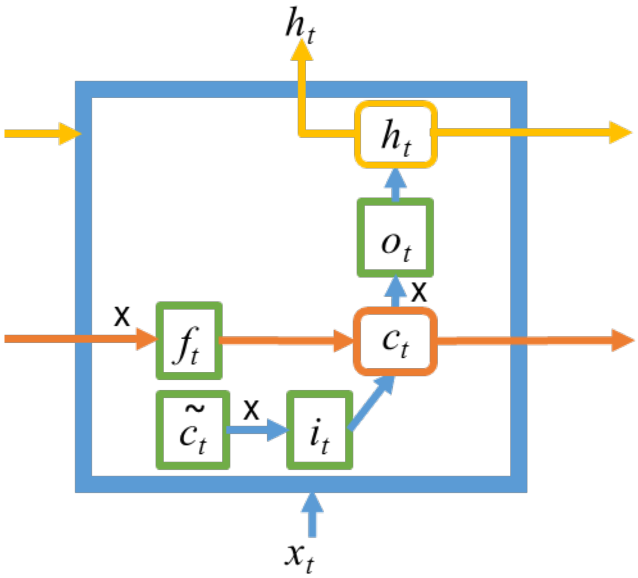

Outline

1. Supervised machine learning for fMRI
2. Deep learning and recurrent neural networks
3. Classification of autism / control from
    - Resting-state fMRI (rsfMRI)
    - rsfMRI + phenotypic data
4. Prediction of autism treatment outcome
5. Final thoughts

Application 1: Identify fMRI biomarkers for autism spectrum disorder (ASD)

1. Train classifier for ASD 2. Extract biomarkers from learned model

Inputs X

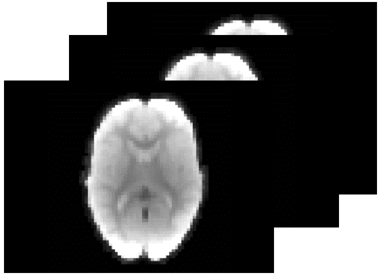

Outputs Y

ASD or TDC

Learning Algorithm

?

Learned Predictive Model

$f : X \rightarrow Y$

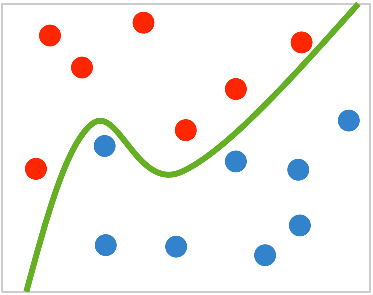

Heterogeneity of ASD makes classification challenging

- Studies using small (N<100), homogeneous datasets:
  - \+ High accuracy (>70%) within sample
  - – Poor generalization
- Studies using large, heterogeneous datasets:
  - \+ Better representation of population
  - – Low accuracy (~60%), likely due to heterogeneity

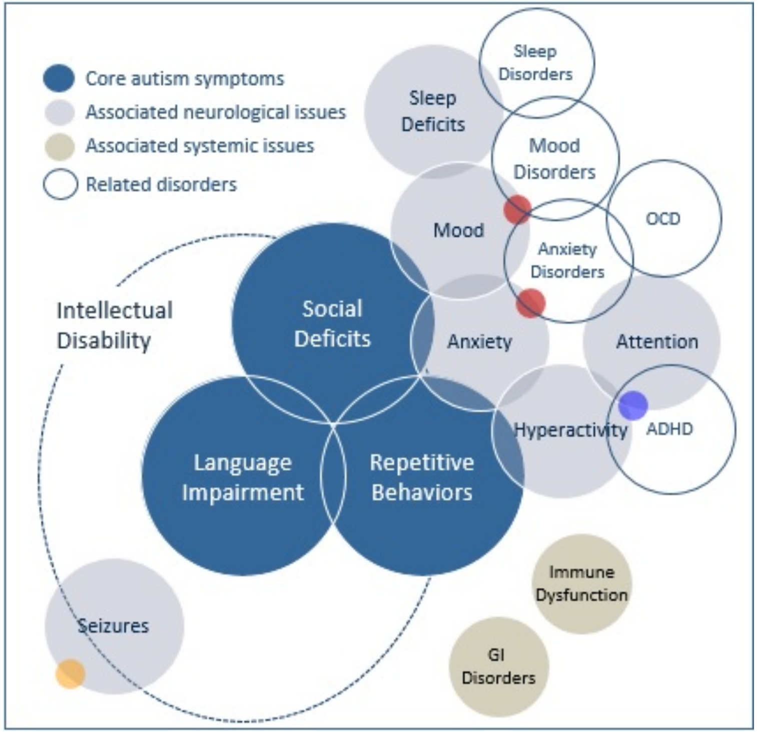

[1]

[1]www.autismspeaks.org

Standard approaches classify ASD/TDC using pairwise rsfMRI connectivity as inputs

Inputs X

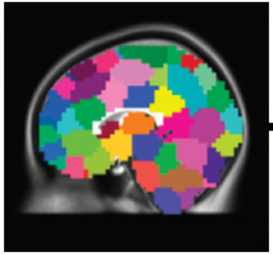

Brain Parcellation

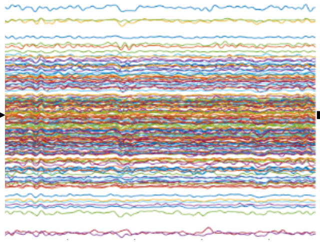

rsfMRI Mean Time-series

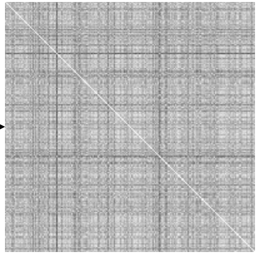

Connectivity Matrix

Outputs Y ASD or TDC

Learning Algorithm

- Support vector machines
- Random Forest
- Logistic Regression

⋮

Learned Predictive Model

$f: X \rightarrow Y$

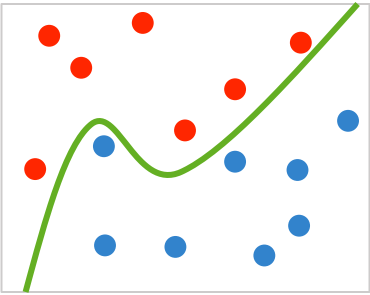

*Biomarkers: Pairwise connections

Proposed approach classifies ASD/TDC using rsfMRI time-series as inputs

Inputs X

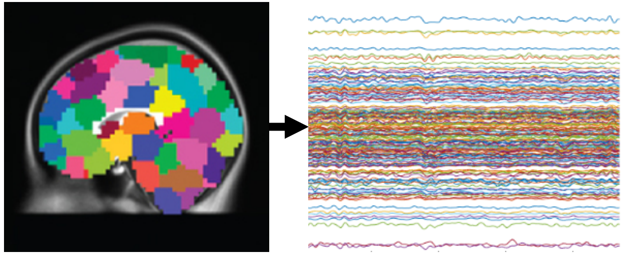

Brain Parcellation

rsfMRI Mean Time-series

Outputs Y

ASD or TDC

Learning Algorithm

- LSTM-based Network

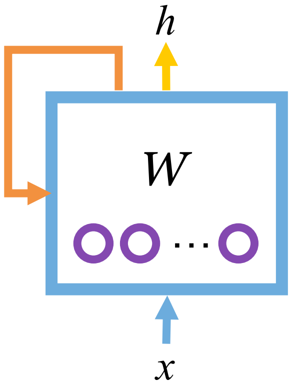

Learned Predictive Model

$f : X \rightarrow Y$

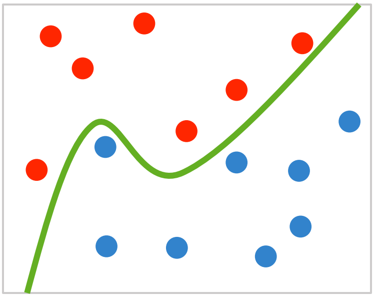

*Biomarkers: Brain regions/networks defined by LSTM nodes

LSTM-based network architecture to classify ASD/TC from rsfMRI time-series

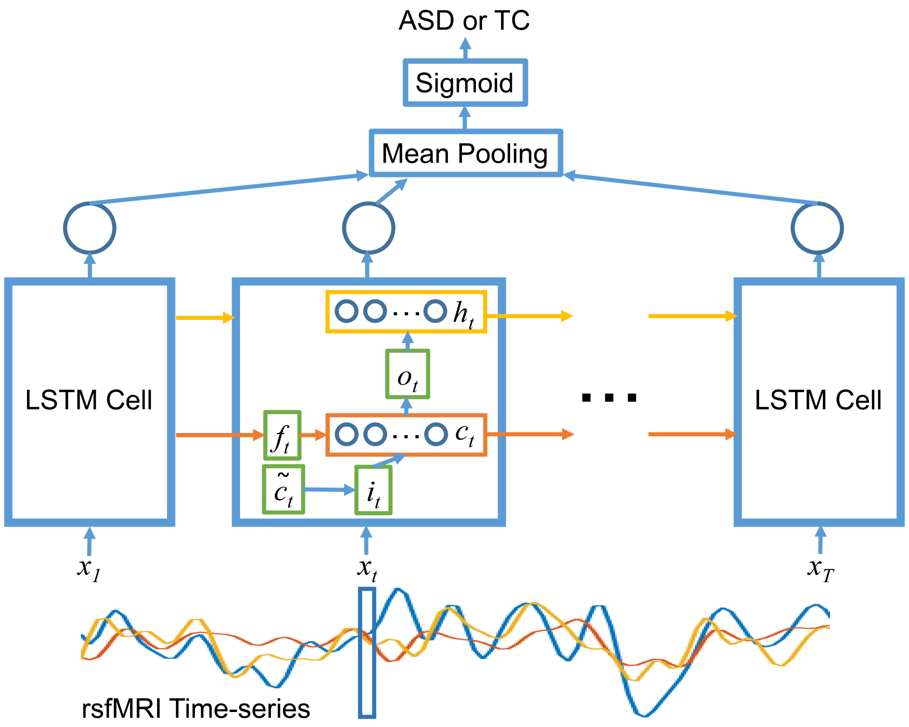

Autism Brain Imaging Data Exchange (ABIDE) Dataset I

- 539 ASD + 573 TDC from 17 international sites
- Neuroimaging (rsfMRI, structural MRI) and phenotypic data
- Preprocessed Connectomes Project
  - Connectome Computation System pipeline → **1100** preprocessed subjects in MNI space
  - Parcellation from Craddock 200 atlas

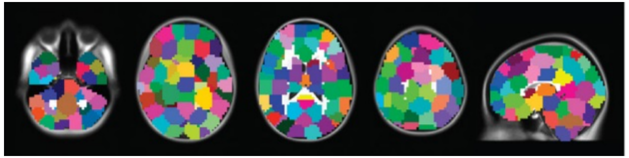
[1]

- ROI mean time-series preprocessing
  - Normalized to percent signal change
  - Resample to 2 second interval

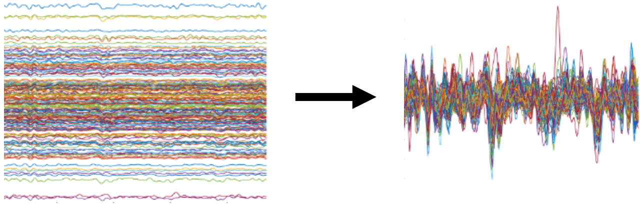

[1]Craddock et al., Nature Methods 2013

Augment data using random crops of the rsfMRI time-series

- 10 random crops / subject → N = 11000 samples

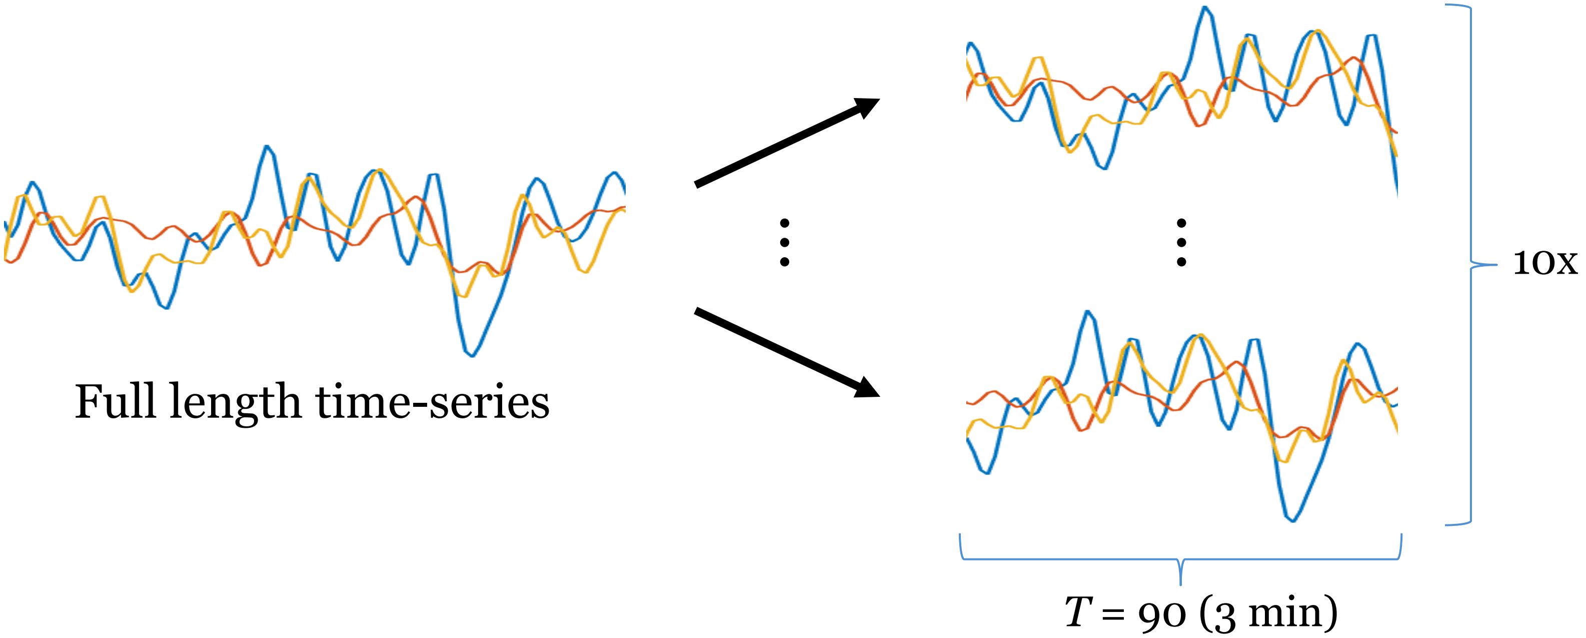

Evaluate algorithm performance using cross-validation

- 10-fold cross-validation, stratified across sites

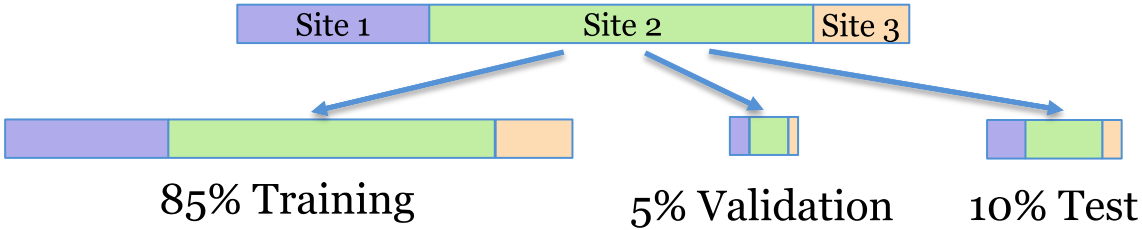

- Accuracy measures
  - Subject accuracy (subject score = average of subject's 10 sample scores)
  - ΔBaseline = (Subject accuracy) – (Accuracy by naïve labelling)

LSTM model outperformed previous studies classifying majority of ABIDE cohort from rsfMRI

Classification Method	Validation Method	# of Subjects	Subject Accuracy (%)	ΔBaseline (%)
Logistic Regression[1]	10CV	178	69.7	19.7
Linear SVM[2]	Train/Val	252	66	16
Linear SVM[3]	10CV	871	66.9 (2.7)	13.2
GLM[4]	LOOCV	964	60.0	6.4
RBF SVM[5]	Train/Val	1111	59.2	7.6
LSTM – no augmentation[6]	10CV	1100	61.4 (4.5)	9.5
LSTM – with augmentation[6]	10CV	1100	**68.5 (5.5)**	**16.6**

[1]Plitt et al., Neuroimage: Clin. 2015. [2]Chen et al., Neuroimage: Clin. 2015. [3]Abraham et al., Neuroimage 2017. [4]Nielsen et al., Front. Hum. Neurosci. 2013. [5]Ghiassian et al., PLOS One 2016. [6]Dvornek et al., MLMI 2017.

Interpret large LSTM model weights as important features for classification

- $W_l(n,r)$ = Weight for ROI r in LSTM node n for layer l

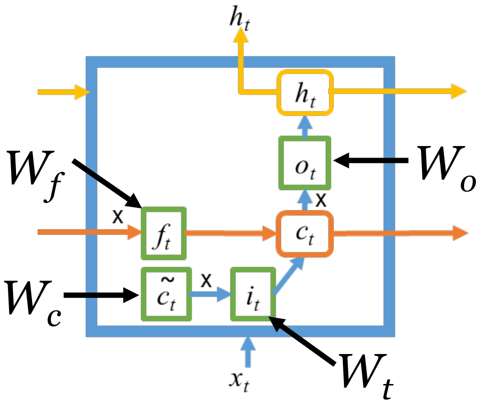

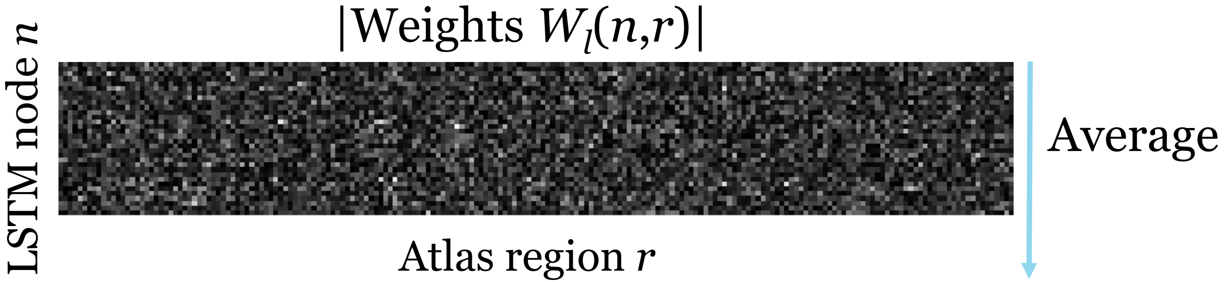

analysis	corr.
posterior superior	0.481
temporal sulcus	0.472
superior temporal	0.402
psts	0.396
temporal	0.388
sts	0.386
⋮	

Neurosynth

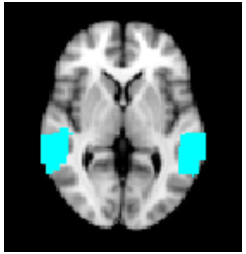

Mask of large average $|W|$

Important ROIs for each LSTM layer map to areas previously implicated in ASD

Input Layer

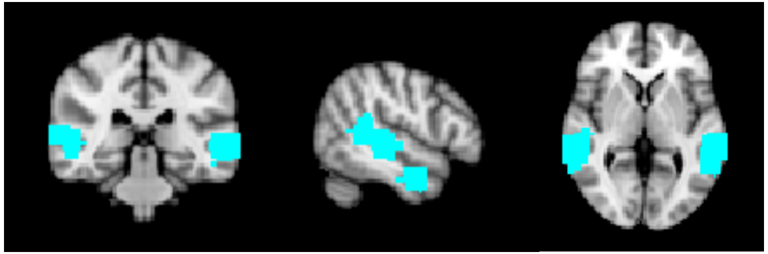

Forget Layer

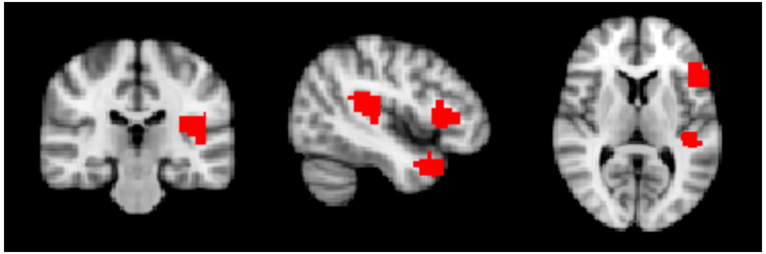

	Input Layer	Forget Layer
Anatomical:	Superior Temporal Sulcus, Middle Temporal Gyrus, Planum Temporale	Inferior Frontal Gyrus, Temporal Pole, Planum Temporale
Functional:	Sentence, Comprehension, Linguistic, Audiovisual, Language	Sentence, Verb, Nouns, Semantically, Sentence Comprehension

Important ROIs for each LSTM layer map to areas previously implicated in ASD

	Cell Layer	Output Layer
	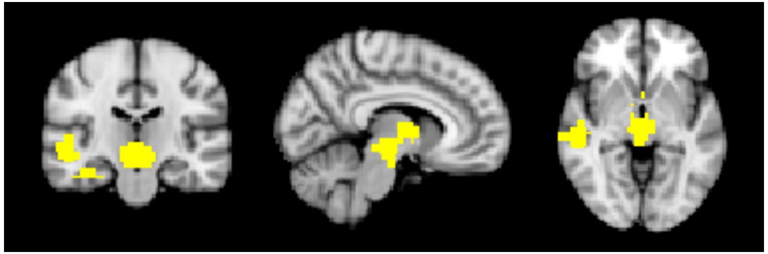	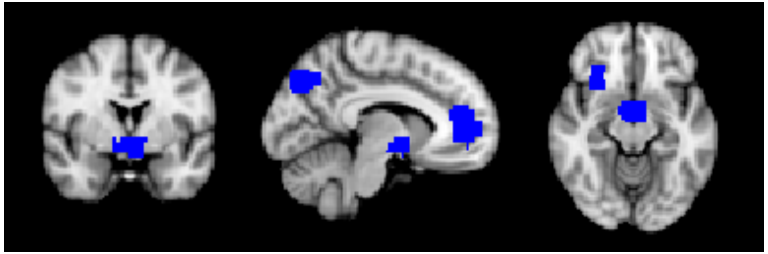
Anatomical:	Midbrain, Thalamus, Superior Temporal Sulcus	Hypothalamus, Inferior Parietal Lobe, Medial Prefrontal Cortex
Functional:	Reward, Speaker, Voice, Audiovisual, Speech	Self, Sexual, Referential, Memory Retrieval, Regulation

Interpret large LSTM model weights as important features for classification

- $W_l(n,r)$ = Weight for ROI r in LSTM node n for layer l

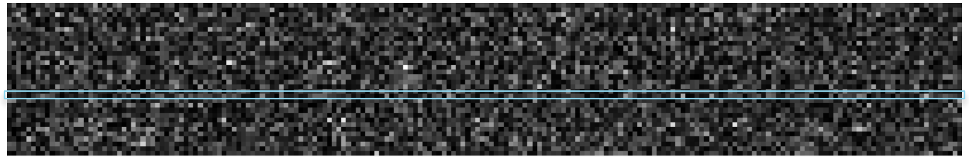

Mask of large $|W|$

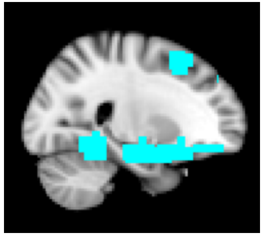

Neurosynth decoding

analysis	corr.
fusiform	0.205
fusiform gyrus	0.179
occipitotemporal	0.171
lateral occipital	0.17
faces	0.167
objects	0.167

⋮

Each LSTM node encodes potential brain network important for ASD/TDC classification

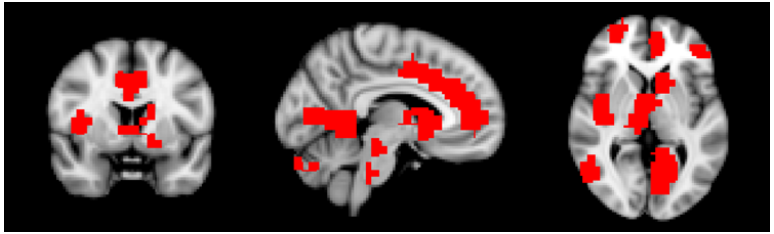

Pain, reward, anticipation, incentive

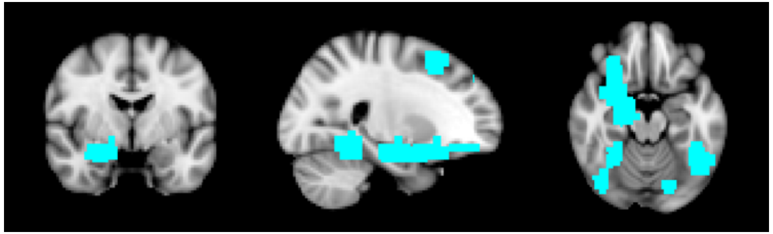

Faces, objects, word form, emotional, visual

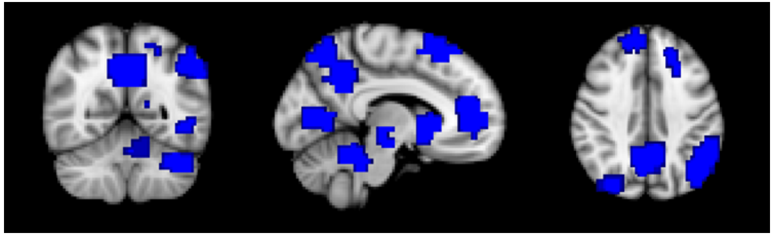

Default mode, reward, listening, mental states

Listening, sounds, theory of mind, social

Outline

1. Supervised machine learning for fMRI
2. Deep learning and recurrent neural networks
3. Classification of autism / control from
   - Resting-state fMRI (rsfMRI)
   - rsfMRI + phenotypic data
4. Prediction of autism treatment outcome
5. Final thoughts

How to incorporate phenotypic information?

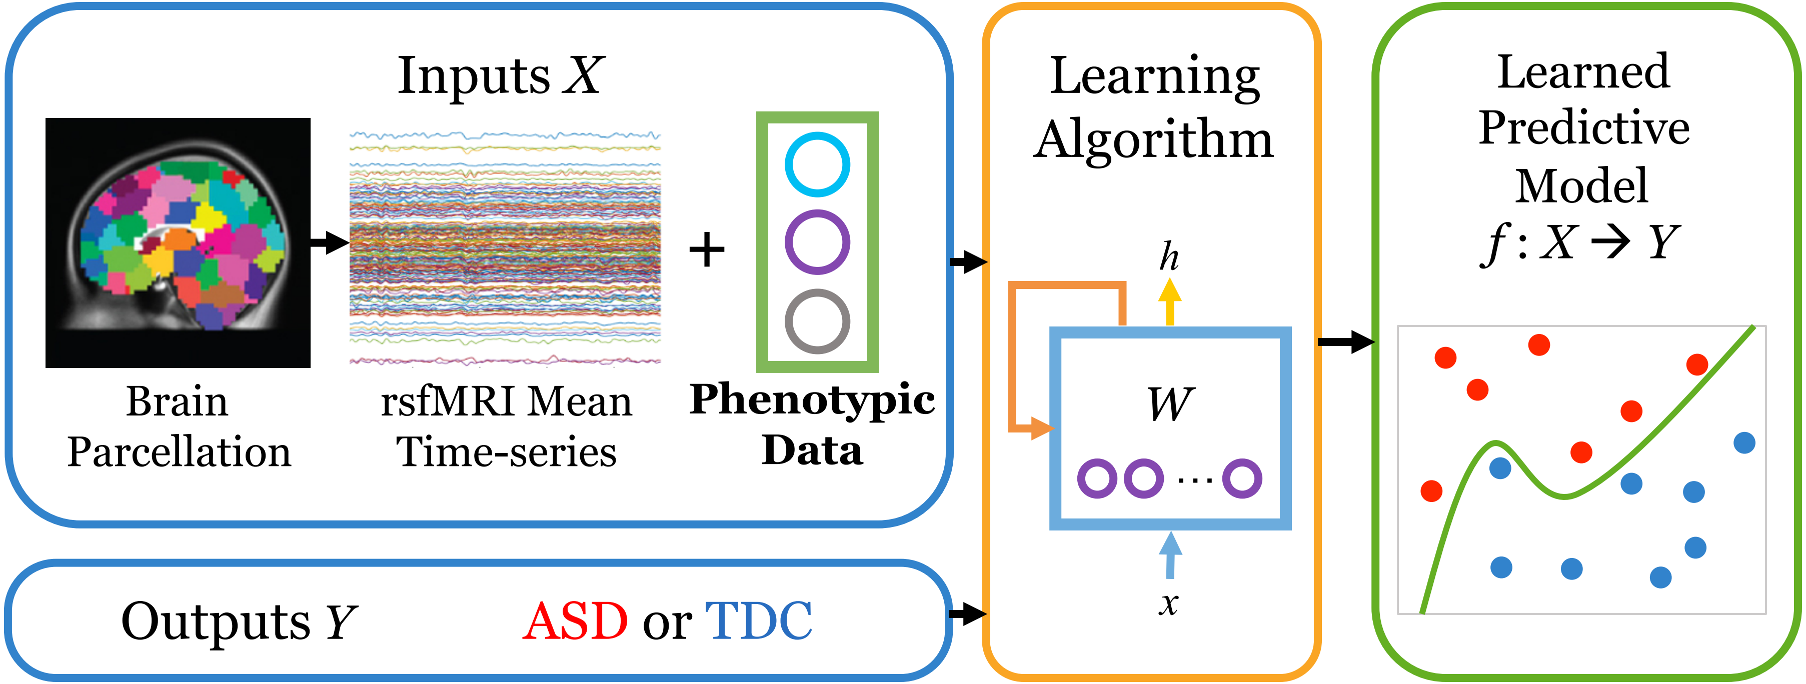

*Biomarkers: Brain networks defined by LSTM nodes

Recall: Baseline LSTM architecture

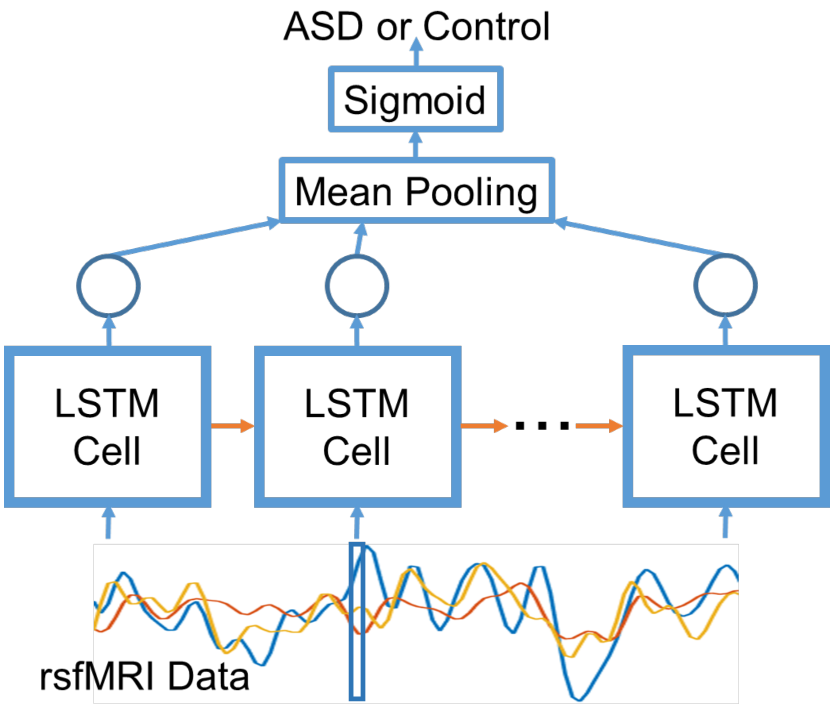

Input phenotypic data with rsfMRI directly into LSTM (Pheno-TS)

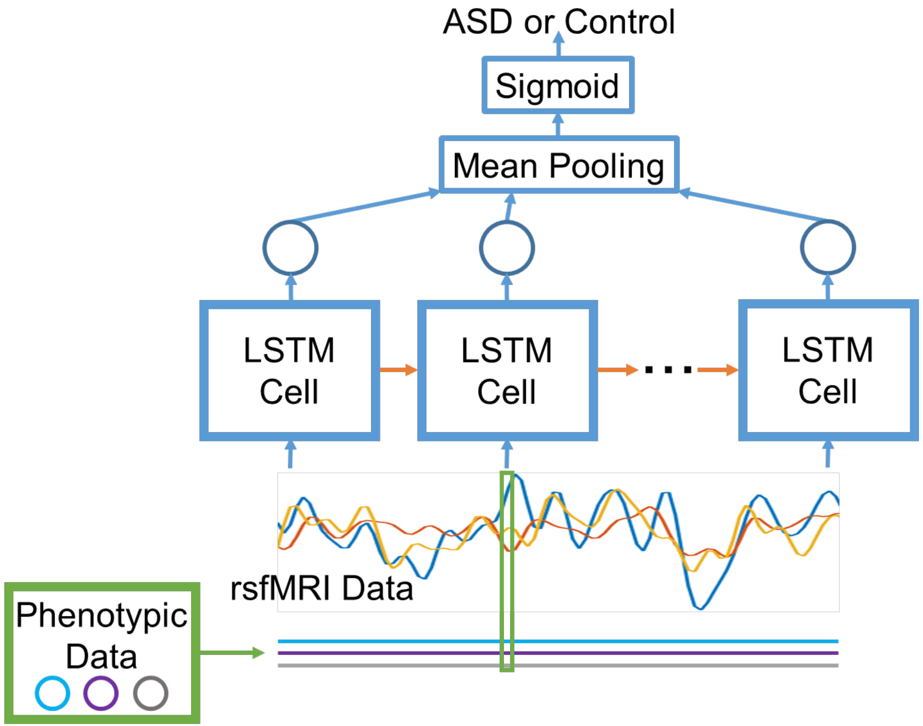

Combine phenotypic data with LSTM encoded outputs (Pheno-LSTM)

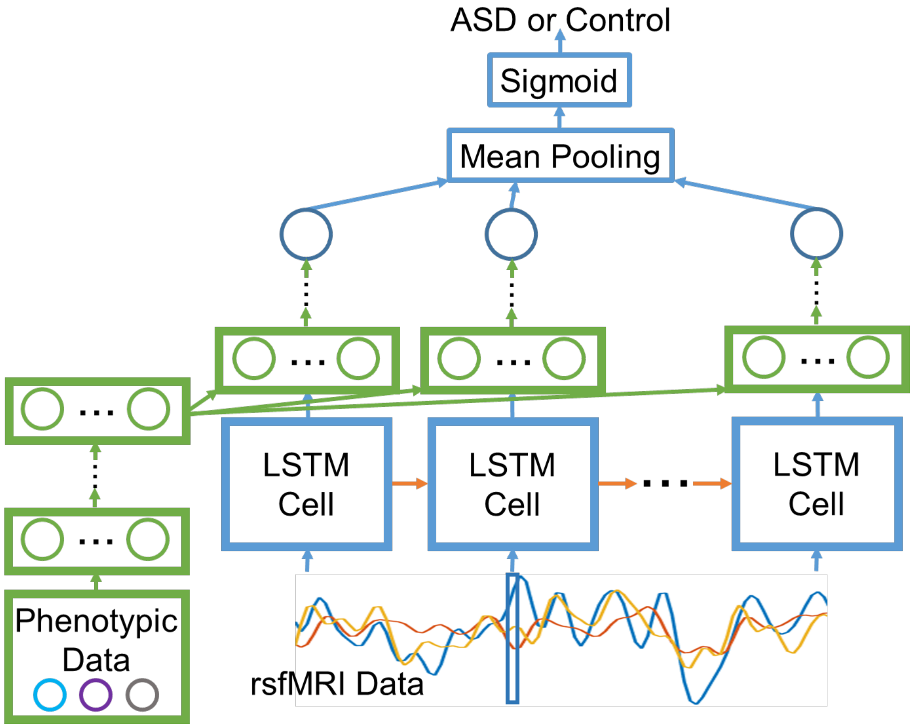

Combine phenotypic data with score from rsfMRI (Pheno-rsfScore)

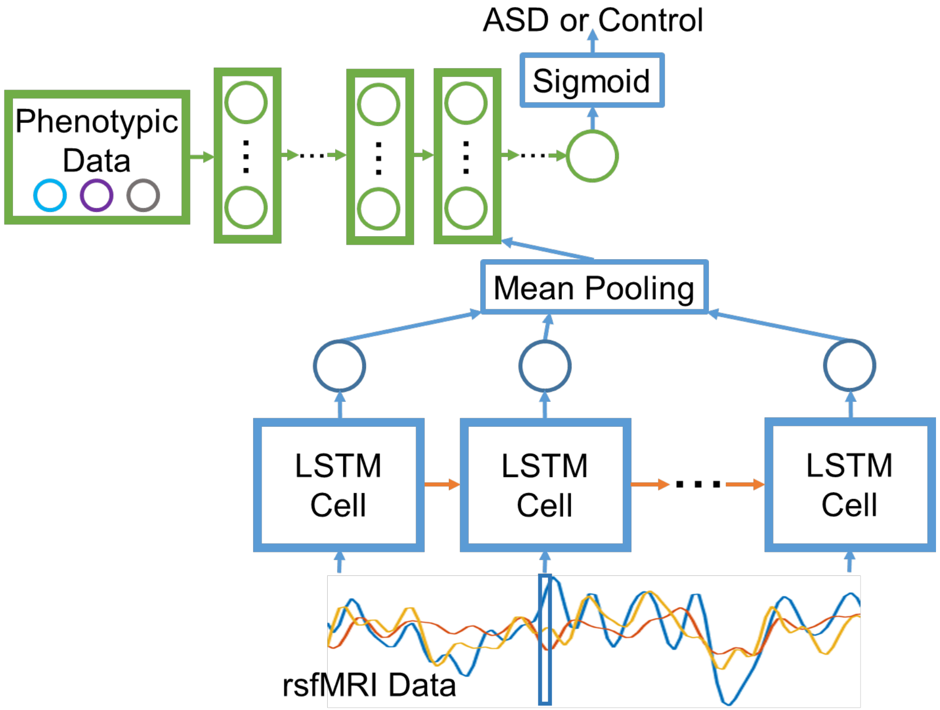

Use phenotypes as auxiliary targets (Pheno-Target)

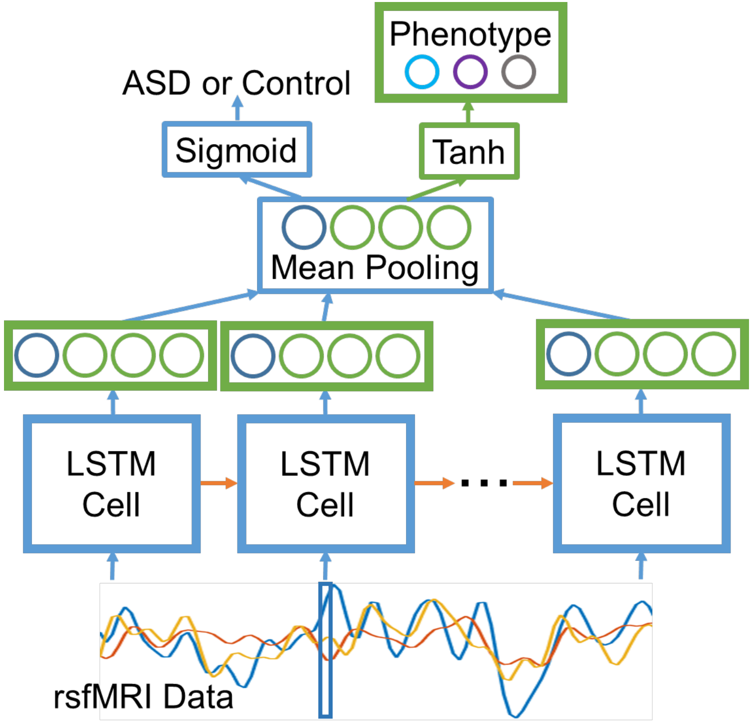

Combining phenotypic data with rsfMRI score improves classification accuracy

Classification Method	Phenotypic Data	# of Subjects	Subject Accuracy (%)	ΔBaseline (%)
Graph CNN[1]	Sex, Site	871	69.5	15.8
GLM[2]	Age, Sex, Hand	964	60.0	6.4
RBF SVM[3]	Age, Sex, Hand, Full IQ, PIQ, VIQ, Site, Eye	1111	65.0	13.4
Pheno-TS[4]	Age, Sex, Hand, Full IQ, Eye	1100	67.0 (3.5)	15.1
Pheno-LSTM[4]			68.2 (4.1)	16.3
Pheno-rsfScore[4]			**70.1 (3.2)**	**18.2**
Pheno-Target[4]			67.2 (4.0)	15.3

[1]Parisot et al., MICCAI 2017. [2]Nielsen et al., Front. Hum. Neurosci. 2013. [3]Ghiassian et al., PLOS One 2016. [4]Dvornek et al., ISBI 2018.

For best model, rsfMRI is most influential and full IQ is most important phenotypic variable

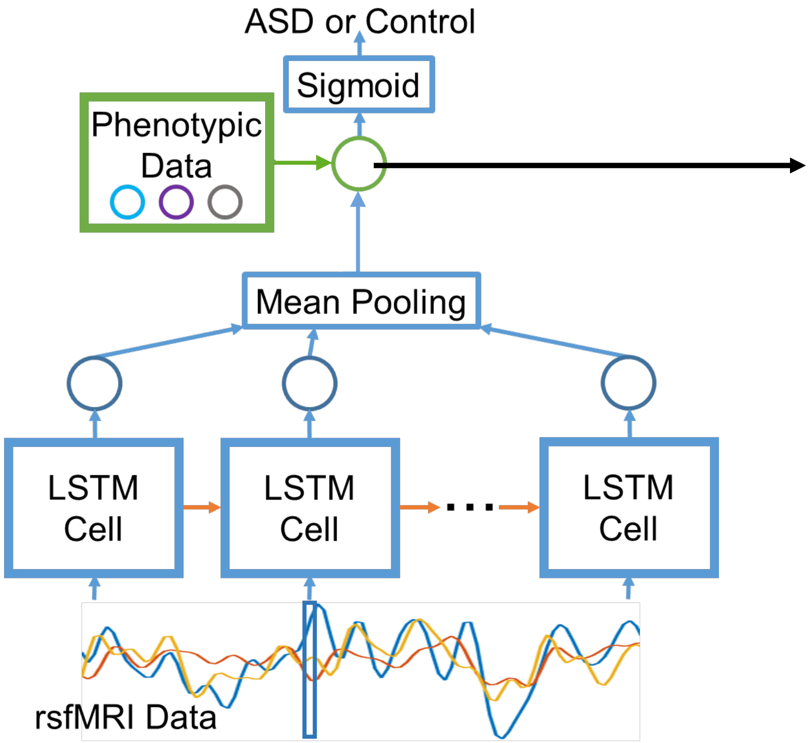

Magnitude of weights for top layer

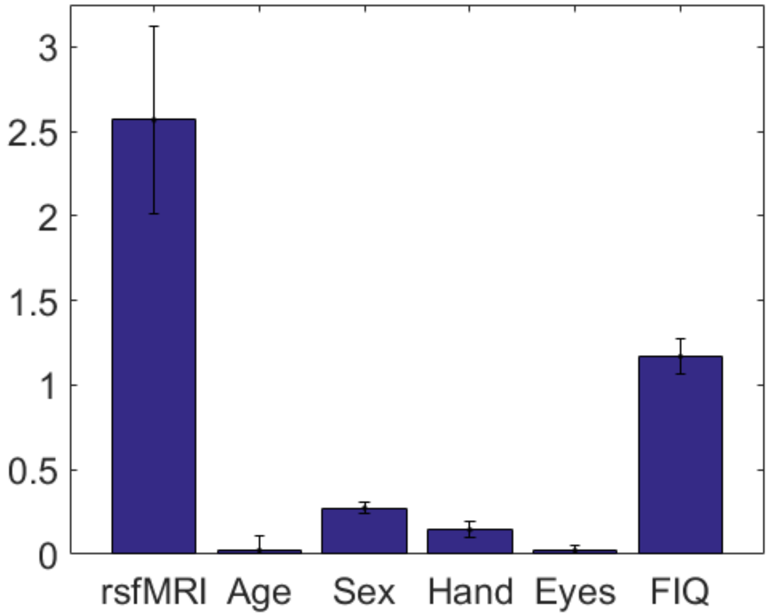

Outline

1. Supervised machine learning for fMRI
2. Deep learning and recurrent neural networks
3. Classification of autism / control from
   - Resting-state fMRI (rsfMRI)
   - rsfMRI + phenotypic data
4. Prediction of autism treatment outcome
5. Final thoughts

Application 2: Predict treatment outcome from baseline fMRI

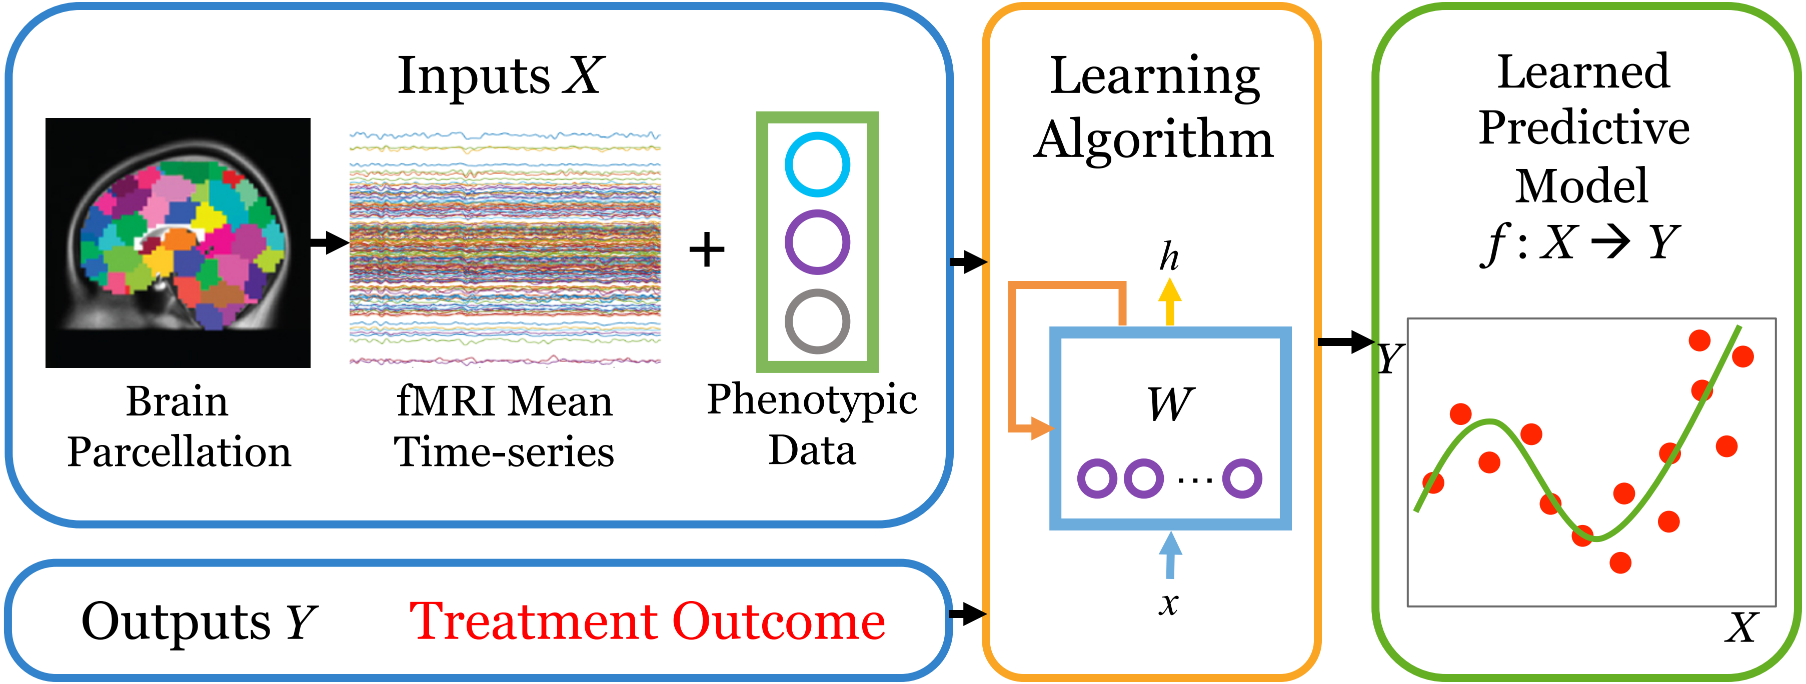

Intensive behavioral therapies may be a promising treatment for ASD

- E.g., Pivotal Response Therapy (PRT)
  - Targets social skills development in play-based format
- Large commitment from patients and families
- Early intervention is crucial

[1]

[1]www.autismspeaks.org

Need for precision medicine to assign correct treatment as early as possible

- Recall: ASD extremely heterogeneous
- No “one size fits all” treatment
- Currently trial and error

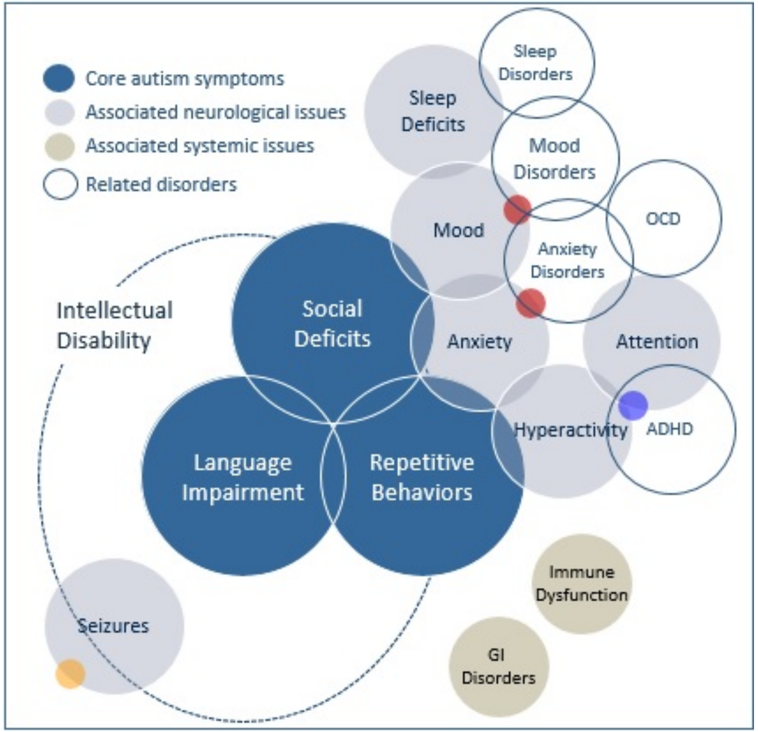

[1]

[1]www.autismspeaks.org

How to learn a robust LSTM model for predicting treatment outcome from small patient dataset?

1. Use task fMRI specific to activating social regions
2. Choose appropriate outcome measure
3. More data augmentation
4. Reduce model complexity
5. Utilize simple phenotypic data to create subject-specific models

Use task fMRI with Biopoint biological motion perception paradigm

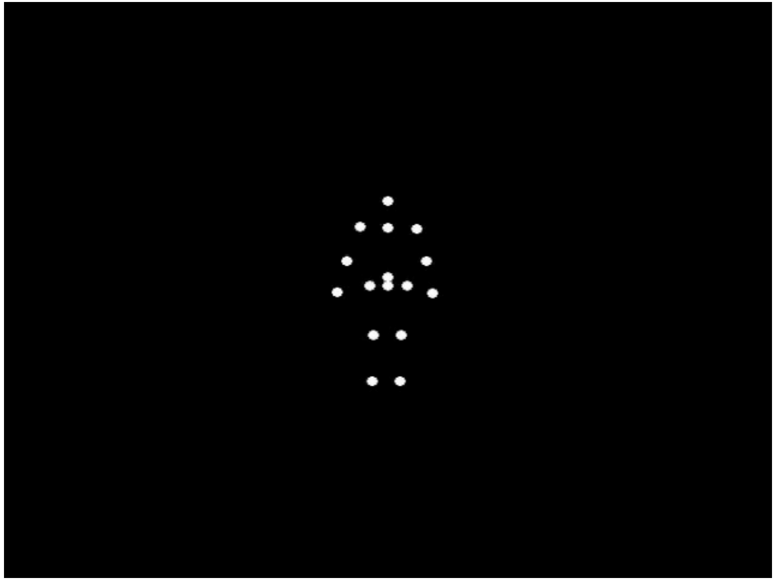

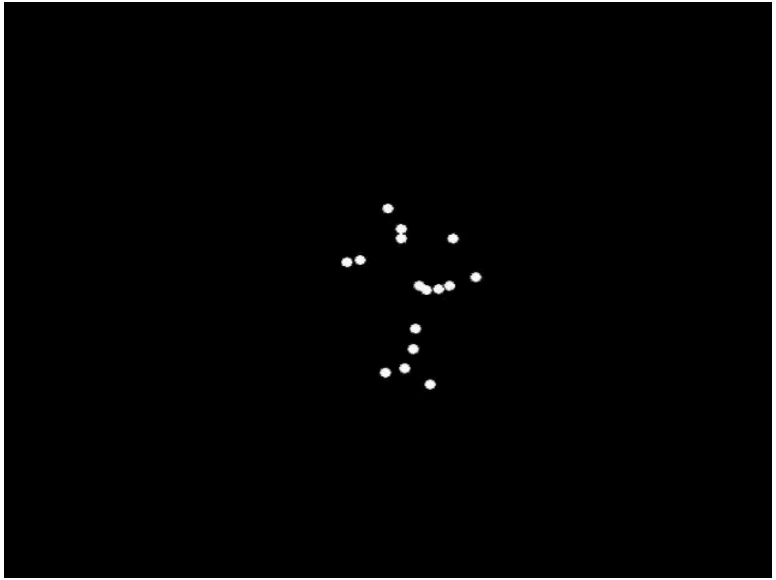

- Revealed differences between ASD/TC in social brain regions[1]

[1]Kaiser et al., PNAS 2010

Measure treatment response using % change in Social Responsiveness Scale (SRS) from baseline

- SRS measures severity of social impairment in autism
- Lower SRS score → Better social function

Baseline SRS	Post-PRT SRS
115	103
71	43
47	52
57	27
106	77
42	40
80	75

⋮

→ Compute % change from baseline

Bootstrap sample ROI voxels to augment small treatment dataset

- For each ROI, randomly sample voxels with replacement

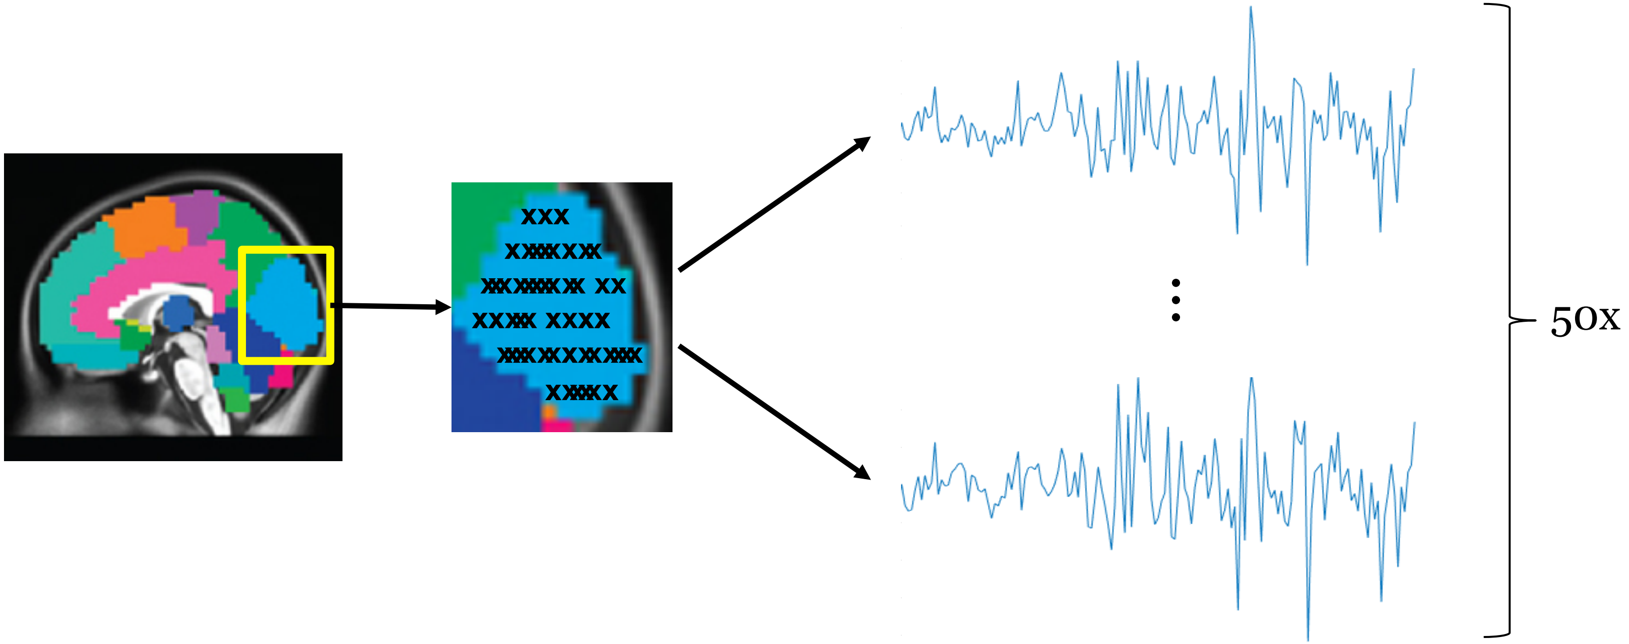

Modify LSTM-based model to predict treatment outcome, conditioned on phenotypic variables

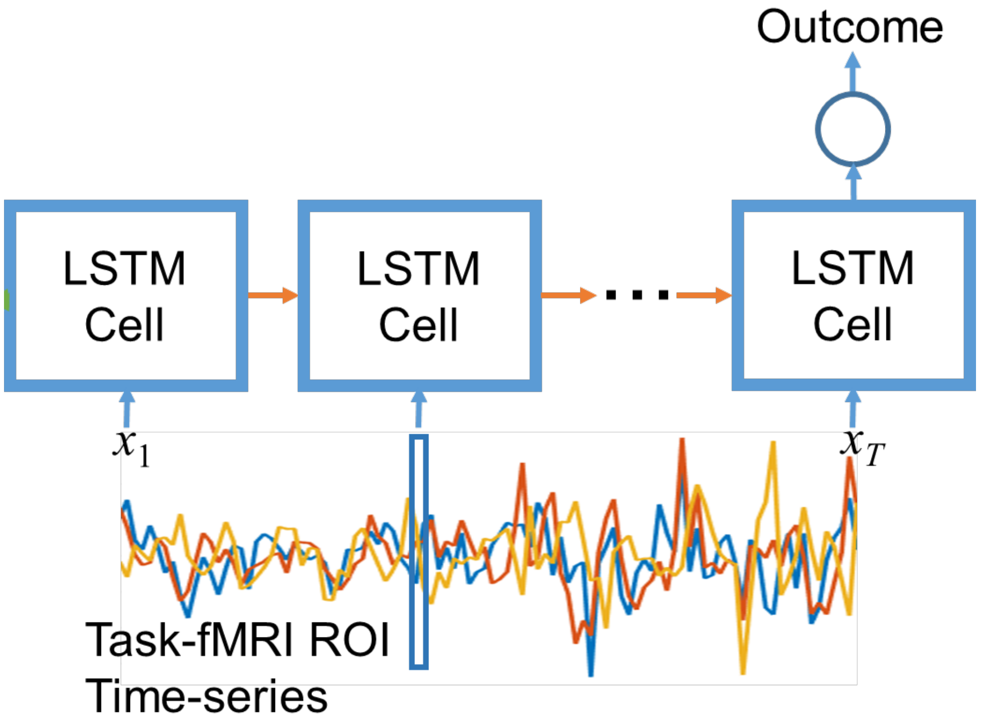

PRT Dataset

- 21 ASD children (6.05 ±1.24 years) underwent 16 weeks PRT
- Collected before treatment:

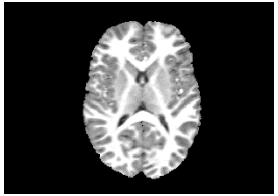

MP-RAGE

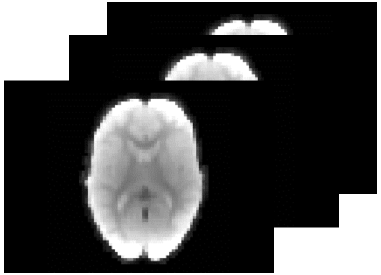

tfMRI with Biopoint

Age
Sex
IQ
SRS

Phenotypic Variables

- Collected after treatment: SRS
- Note: > 2400 hours to collect data

Data Preprocessing

- Standard fMRI preprocessing pipeline[1]
- Registered fMRI to MNI space
- Parcellation using AAL atlas (90 cerebral ROIs)

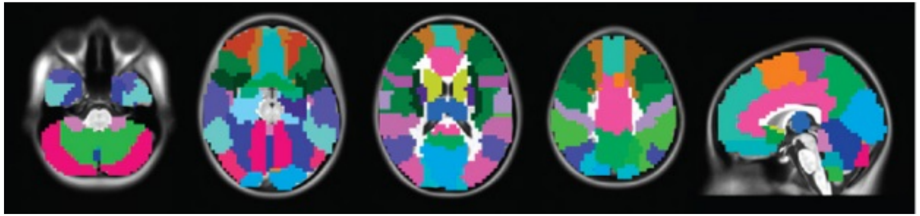

[2]

- Standardized time-series (subtract mean, divide by standard dev.)
- Phenotypic variables normalized to [-1,1]

[1]Pruim et al., Neuroimage 2015. [2]Craddock et al., Nature Methods 2013.

Bootstrap data augmentation and subject-specific LSTM initialization significantly improve prediction

Original Dataset

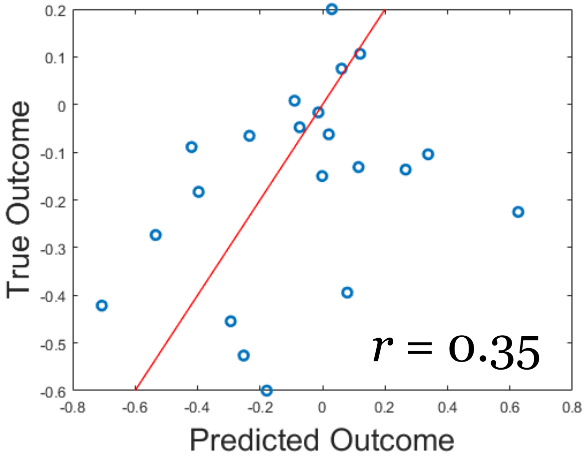

Bootstrap Augmentation

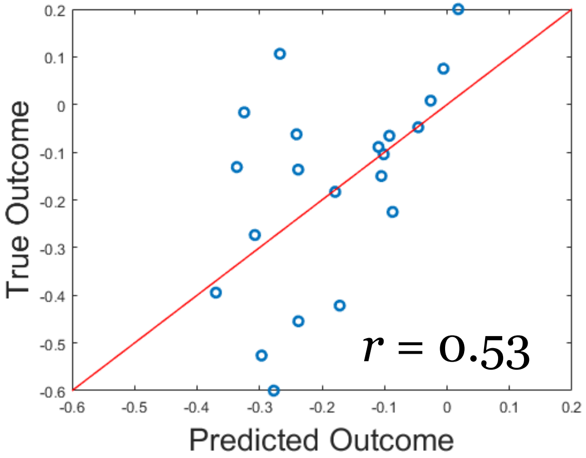

Augmentation + Phenotypes

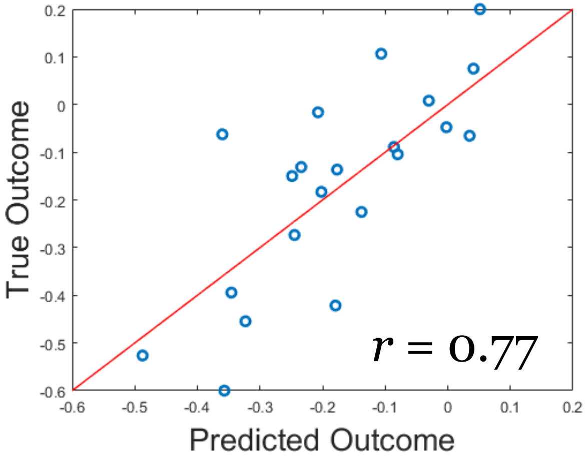

True = Predicted

*Significantly better than original

*Significantly better than augmentation

Outline

1. Supervised machine learning for fMRI
2. Deep learning and recurrent neural networks
3. Classification of autism / control from
    - Resting-state fMRI (rsfMRI)
    - rsfMRI + phenotypic data
4. Prediction of autism treatment outcome
5. Final thoughts

Deep learning with LSTMs can successfully learn predictive models from fMRI

- Higher classification accuracy (68.5%) of ABIDE cohort

Inputs X

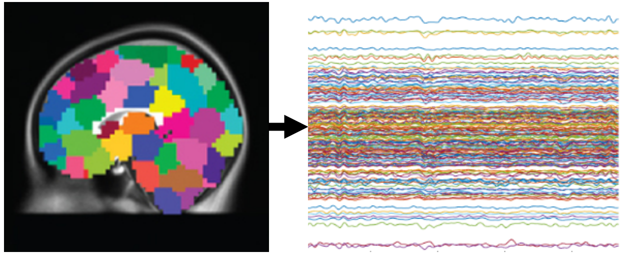

Brain Parcellation

rsfMRI Mean Time-series

Outputs Y ASD or TDC

Learning Algorithm

- LSTM-based Network

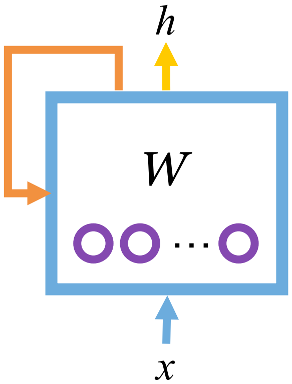

Learned Predictive Model
$f: X \rightarrow Y$

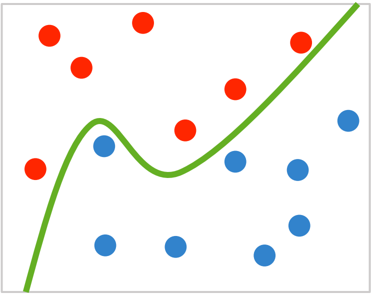

Deep learning with LSTMs can successfully learn predictive models from fMRI + phenotypic data

- Highest classification accuracy (70.1%) on largest subset of ABIDE

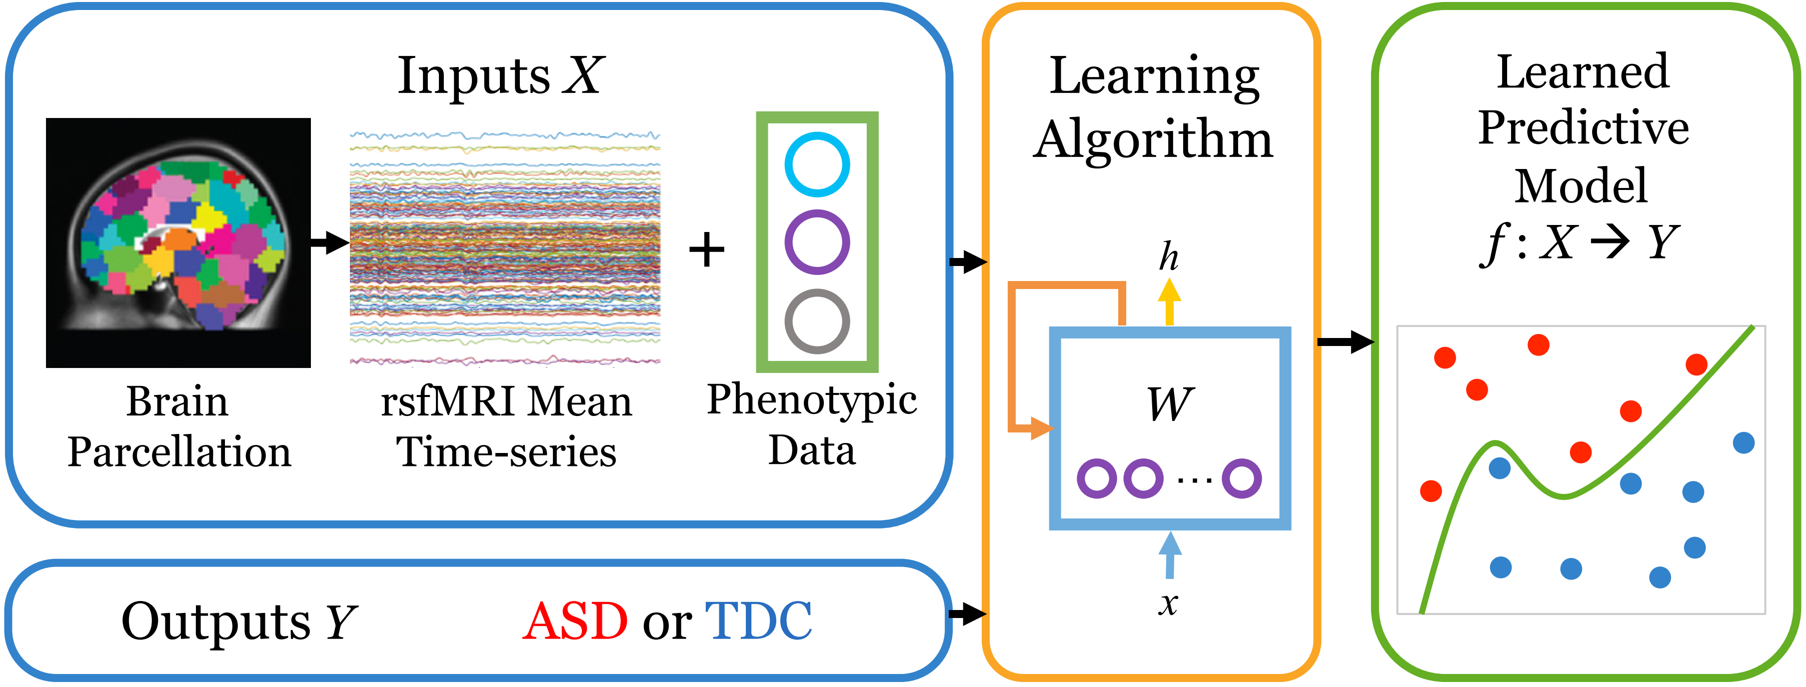

Deep learning with LSTMs can successfully learn predictive models from fMRI + phenotypic data

- High correlation (r = 0.77) between true/predicted response to PRT

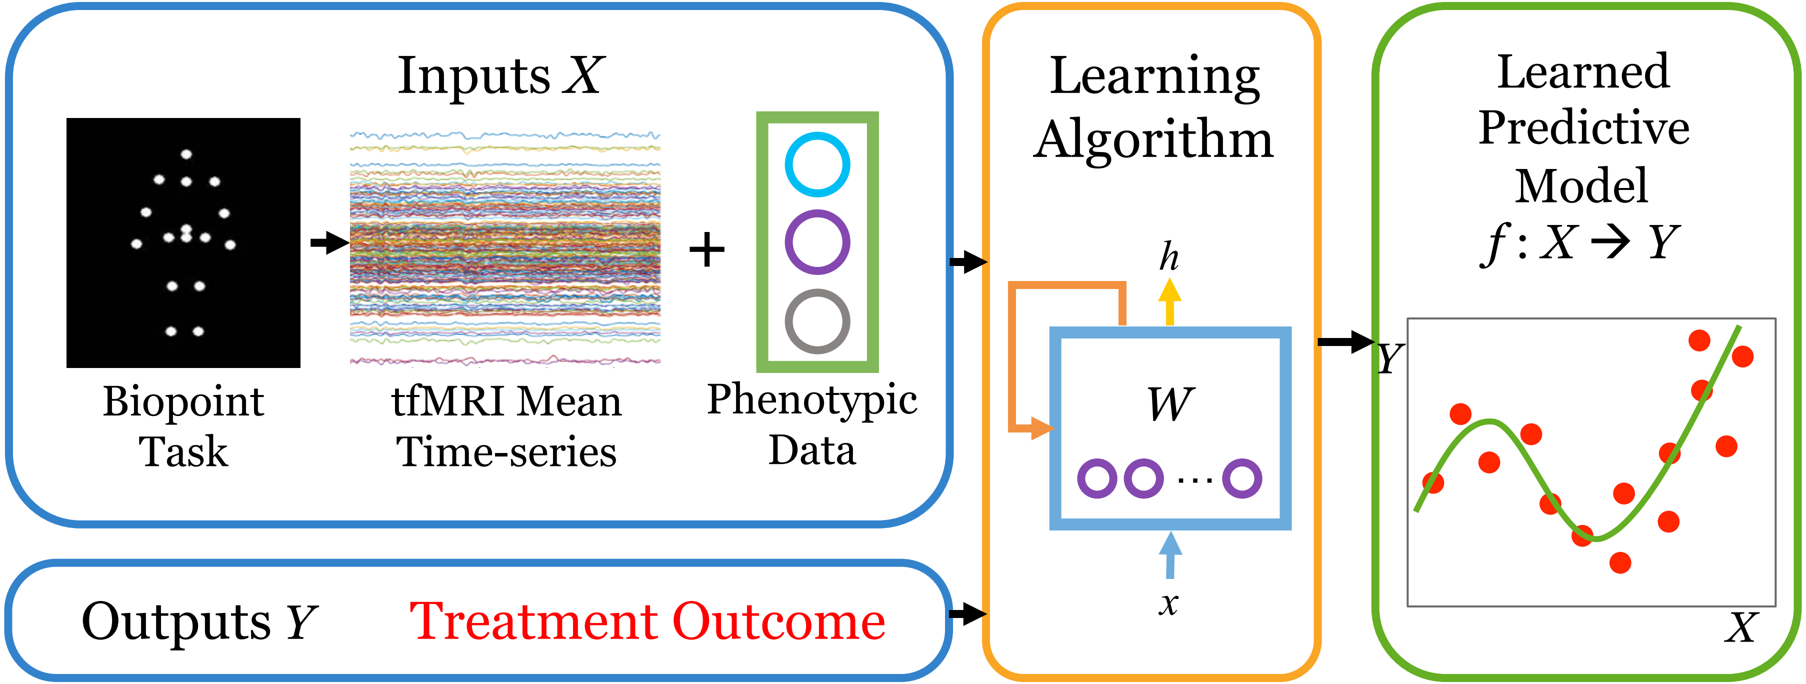

Deep networks with LSTMs are interpretable

- LSTM model weights may define biomarkers
- Extracted regions previously implicated in ASD

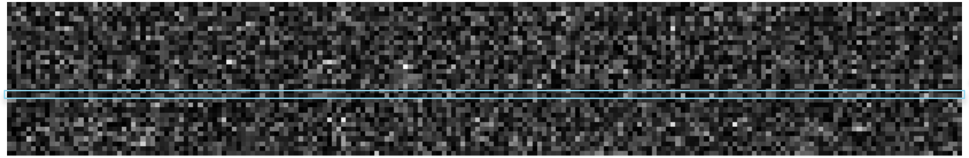

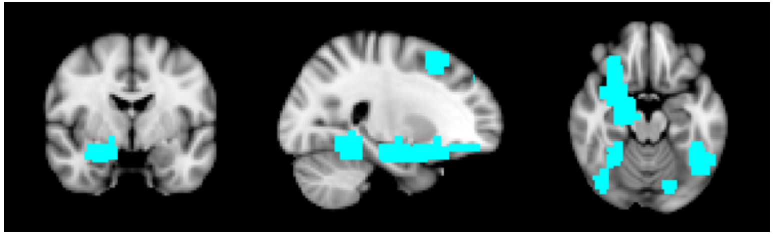

- Depending on model, importance of phenotype can be understood

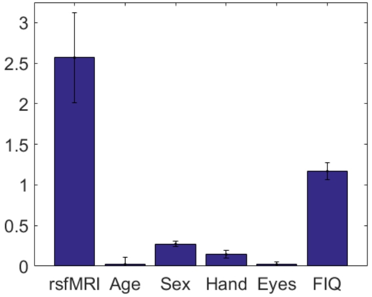

Future directions: New methodology

- Incorporate other types of imaging data into single neural network
  - Structural MRI, Diffusion MRI, …
- Rigorous method for network interpretation/biomarker extraction
- Normalize fMRI data across different sites (domain adaptation)

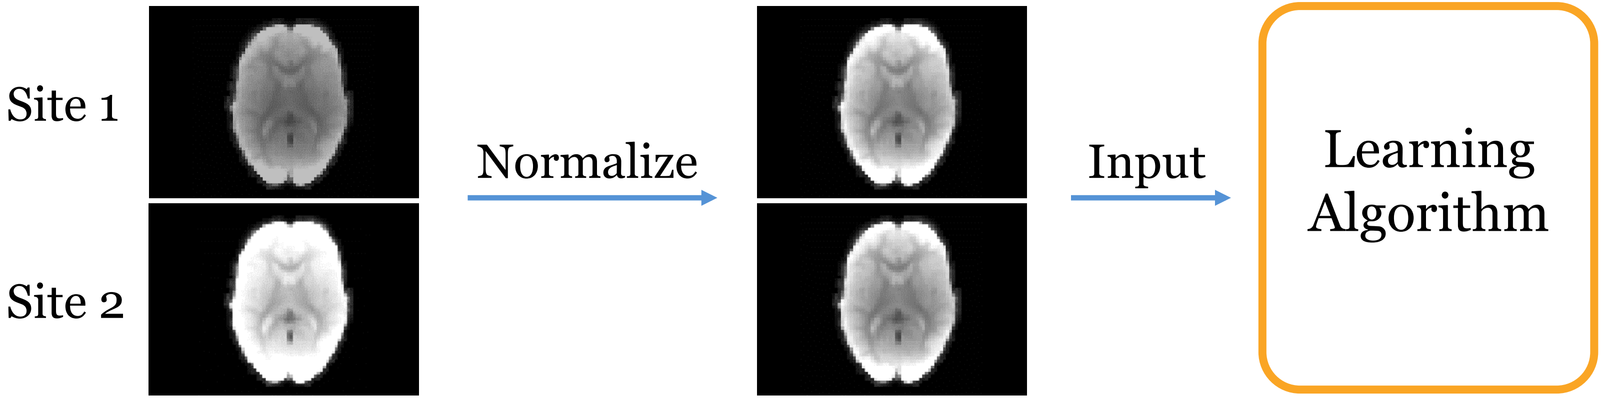

Future directions: New applications

- Other applications in ASD
  - Predict behavioral measures (e.g., SRS, ADOS) and assess biomarkers
  - Predict outcomes for other therapies
- Other neurological diseases/disorders
- Other large neuroimaging datasets, e.g., HCP
- Other neuroimaging modalities
  - EEG
  - MEG
  - PET

In summary….

- Why supervised machine learning for fMRI analysis?
  - Finds *predictive* models that generalize to new data by learning individualized patterns
- What is / why use deep learning with recurrent neural networks?
  - Nice model for temporal data that showed improved accuracy
- How can we apply recurrent neural networks to fMRI analysis?
  - Disease/control classification, biomarker identification, treatment prediction, …

Thank you!

- NIH Grants T32 MH18268 and R01 NS035193
- Duncan Lab
  - Xiaoxiao Li
  - Juntang Zhuang
  - Allen Lu

    ⋮
- Child Study Center (past and present)
  - Pamela Ventola
  - Kevin Pelphrey
  - Daniel Yang

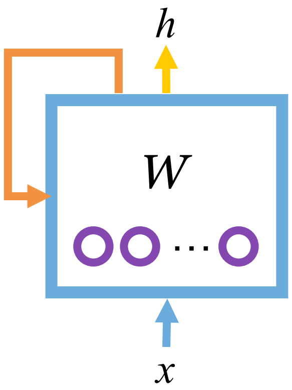

- Contact: nicha.dvornek@yale.edu