Homework: CNN

Neural Network (Q1)

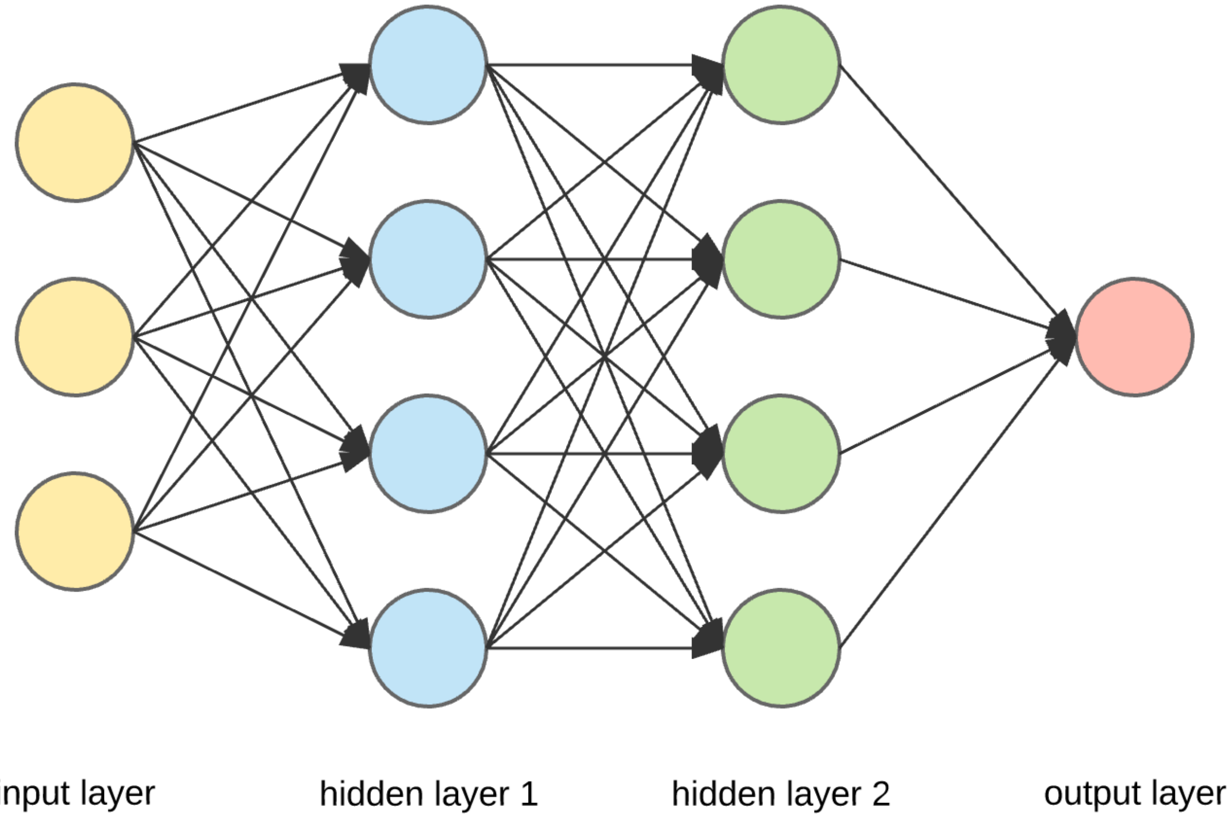

Convolutional Neural Network (Q2)

Conv

Pool

FC

Cross Entropy

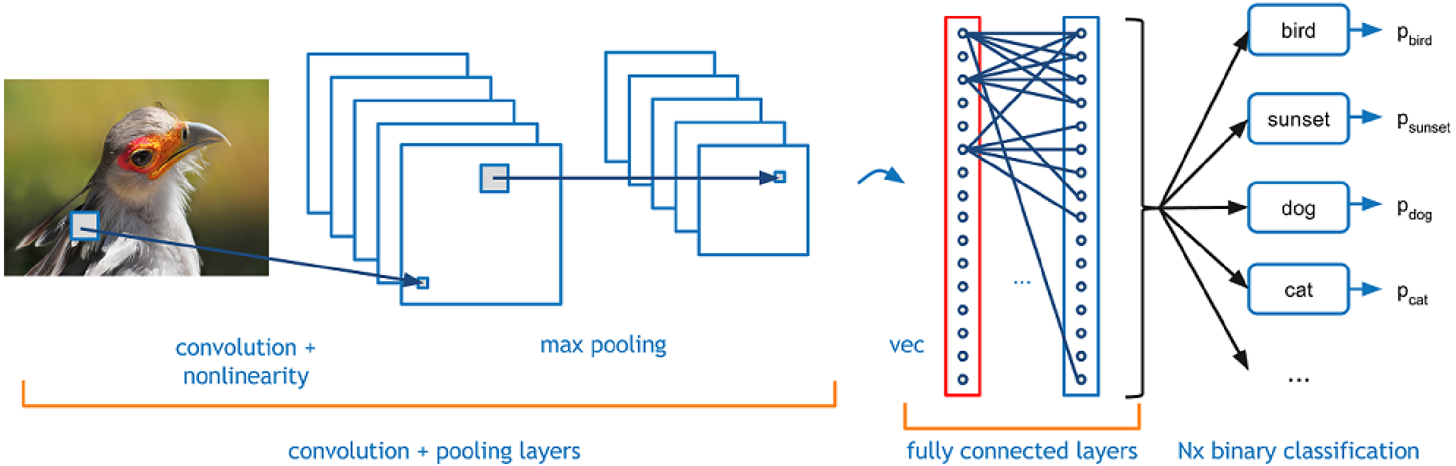

Yellow or Blue?

Color Normalization (Q3)

Deep Convolutional Neural Network (Q4)

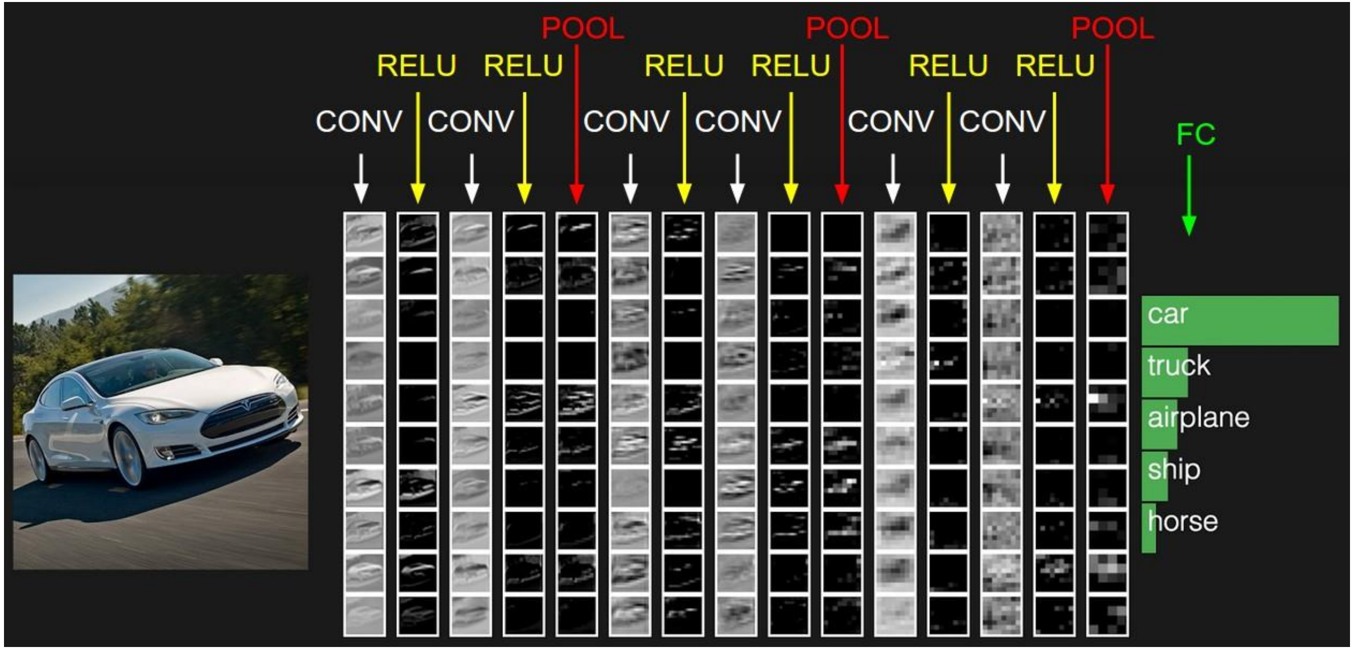

Make the Design More Flexible

Input:

[8, 16, 32, "pool"]

Layer	Output Size	Output Channels
Input	30 x 30	3
Conv	28 x 28	8
ReLU	28 x 28	8
Conv	26 x 26	16
ReLU	26 x 26	16
Conv	24 x 24	32
ReLU	24 x 24	32
Max Pool	12 x 12	32
Linear	5	

Data Augmentation (Q5)

Random Affine Transformation

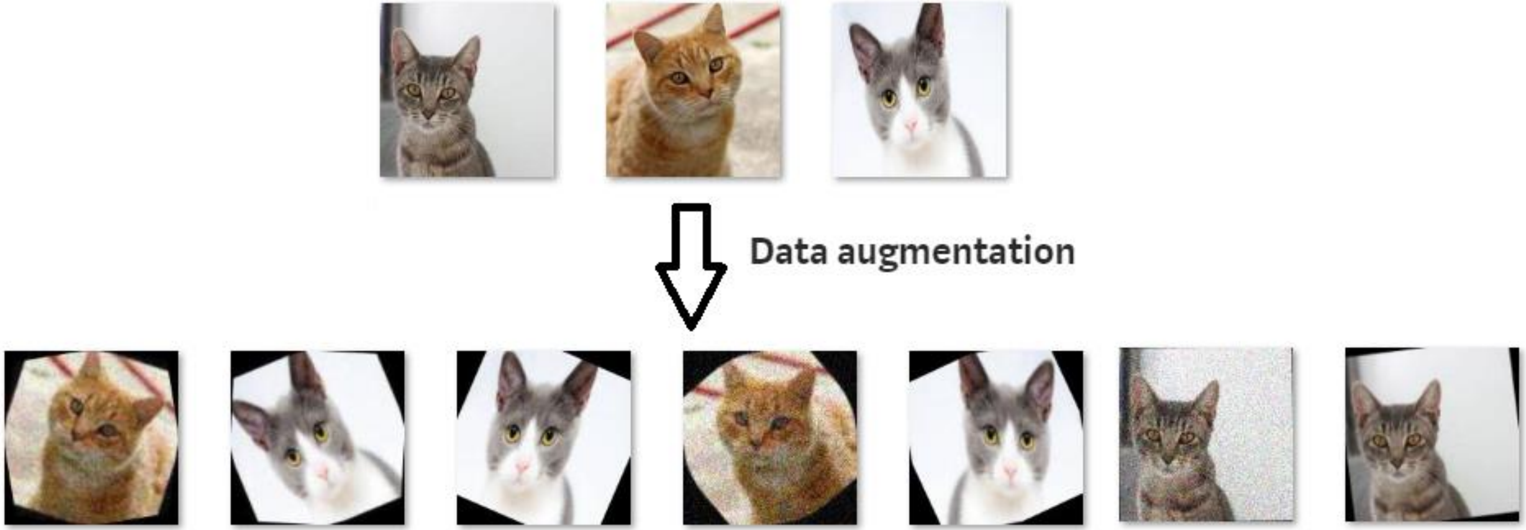

Low-Resource Neural Adaptation

Beibin Li
beibin@uw.edu
Feb/2021

PAUL G. ALLEN SCHOOL

Traditional Machine Learning

Define the Problem → Collect Data → Train Model → Evaluate → Deploy

Breast Cancer

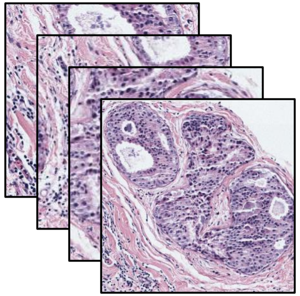

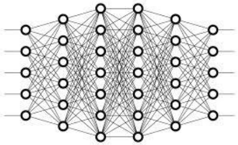

Challenges and Opportunities

During Training

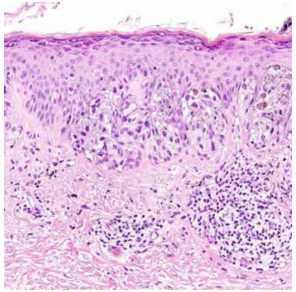

Melanoma

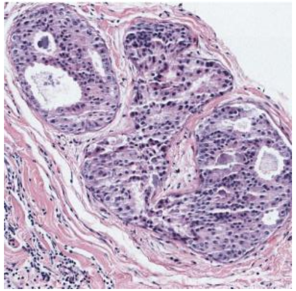

Breast Cancer

During Deployment

Adults' Expression

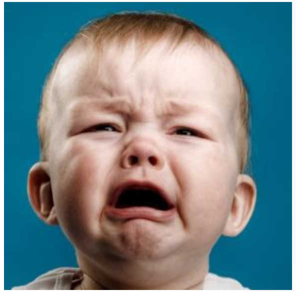

Children's Expression

Adaptation in Humans: Fast and Good

During Learning

Skate

Snowboard

During Using

Drive car

Drive U-Haul Truck

How about Machines?

Can we do **better** and improve **machine**'s adaptation?

Adaptation

Learning

Utilize Prior Knowledge

Practice Different Tasks

Learn from Failure

Inference

Recognize Novel Samples

Estimate Confidence

Adjust Behavior

Adapting Knowledge

Big Data

14+ Million Images

WikiText-103: 100+ Million Tokens

Foundation Models

(Large, Powerful)
(Centralized, Federated)

Small Datasets

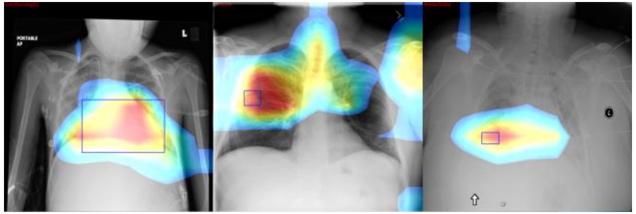

Medical Imaging

Private Documents

Adapting Shifts

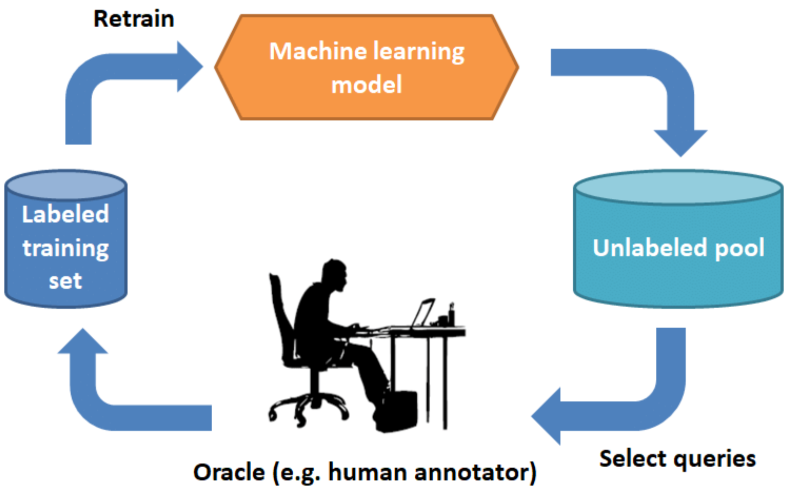

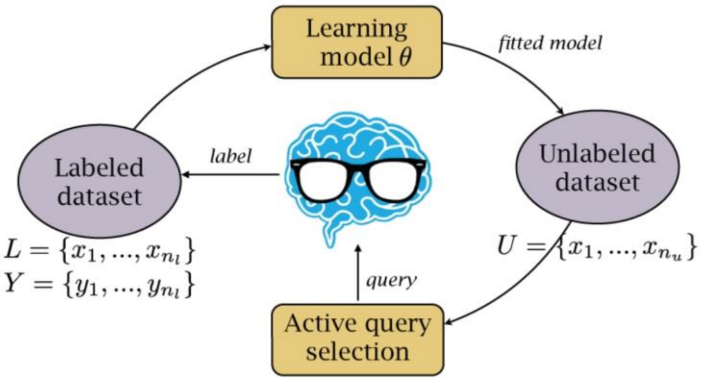

Prior Studies

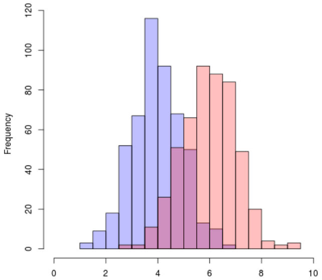

Stats Analysis

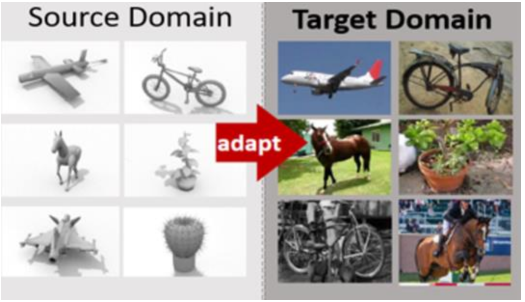

Domain Adaptation

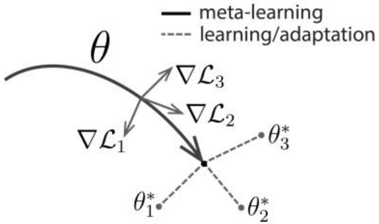

Meta Learning

Self-Supervision

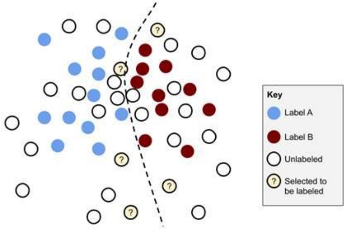

Active Learning

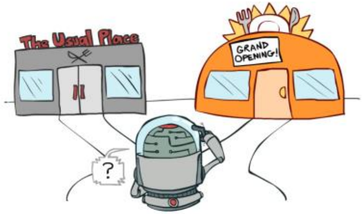

$$R(T) \leq \beta n T + \gamma T + (1+\beta)\eta n T + \frac{\log(n\delta^{-1})}{\beta} + \frac{\log(n)}{\eta}$$

Exploration v.s. Exploitation

They are Powerful, But…

Unified Framework in Human Brain!

Prefrontal Cortex

Executive Functioning Skills
- Shifting
- Inhibitory
- Short-Term Memory

Confidence in Problem-Solving
Language Generation

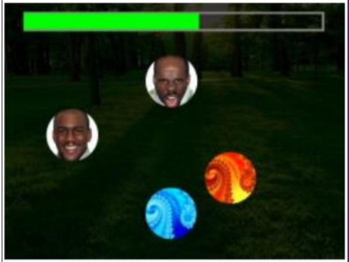

Social Influences on Executive Functioning in Autism: Design of a Mobile Gaming Platform. Li, B., et al. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 443) (ACM SIGCHI 2018).

Memory Deficit in Patients with Temporal Lobe Epilepsy: Evidence from Eye Tracking Technology Zhu, G., et al. Frontiers in Neuroscience 2021

Translate DL for Real World

Bringing state-of-the-art adaptation techniques into real-world applications for private entities (e.g., due to HIPAA, privacy, labeling cost, etc.).

Motivation

We need a **Unified** framework for these various ML tasks

Build **in-house** private foundation models for **low-resource** data

Method

An unified pipeline to adapt shifting in low-resource learning

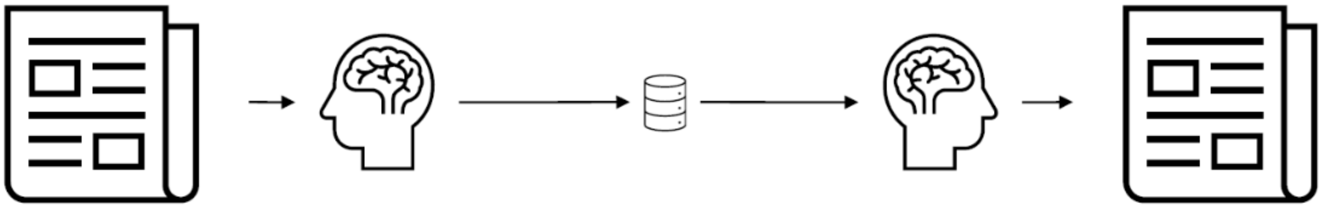

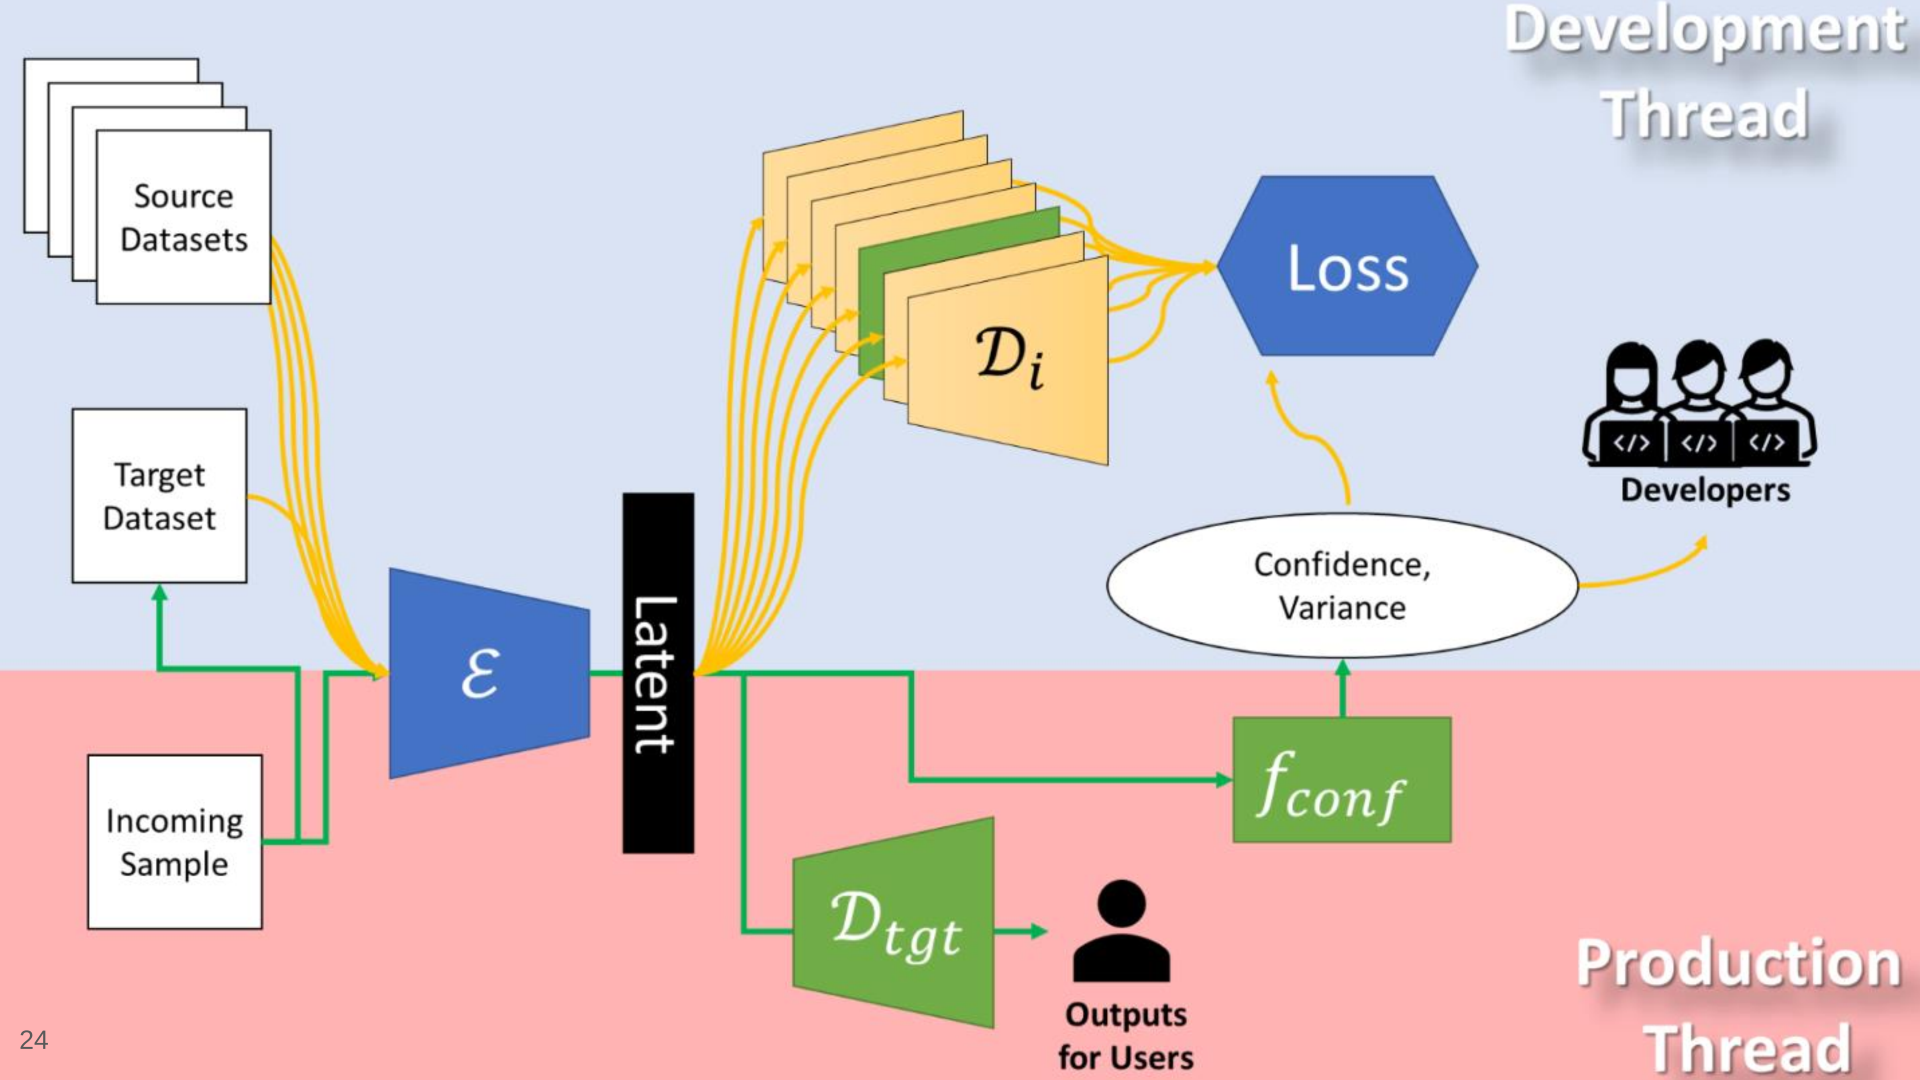

Development Thread
Source Datasets
Target Dataset
Incoming Sample
$\mathcal{E}$
Latent
$\mathcal{D}_i$
Loss
Developers
Confidence, Variance
f_{conf}
$\mathcal{D}_{tgt}$
Outputs for Users
Production Thread

Encoder

- Residual Option
- Variational Bayes Option
- Dropout Mechanism

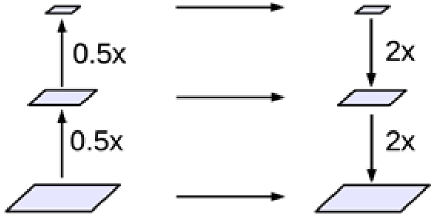

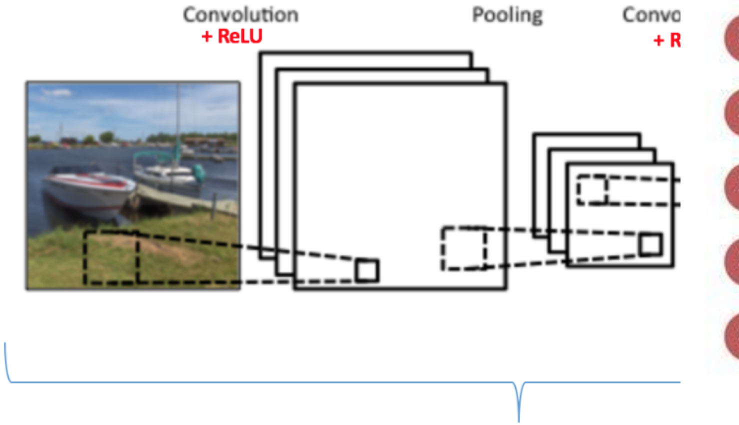

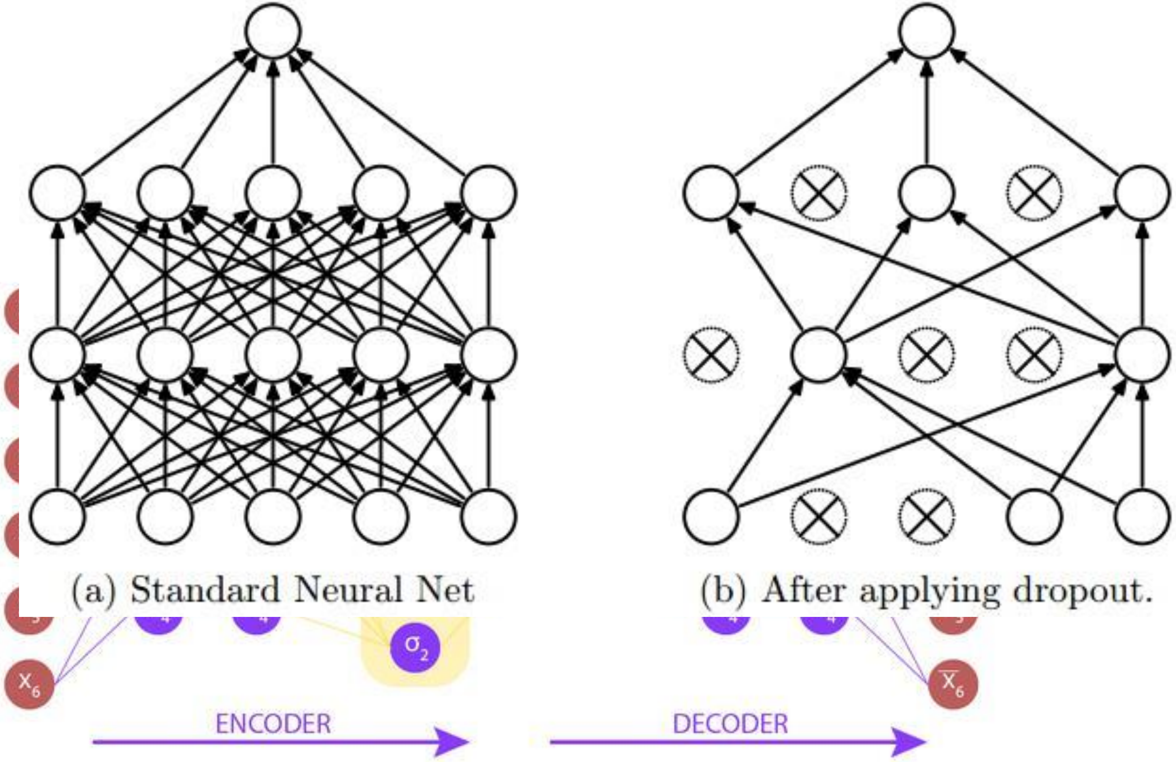

(a) Standard Neural Net (b) After applying dropout.

Decoders and Tasks

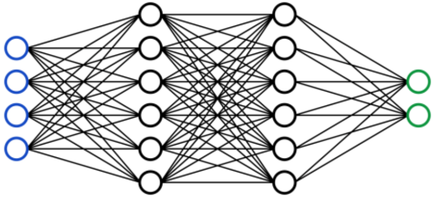

Classify, Regress

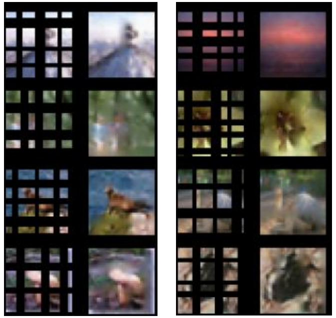

Reconst, Recover

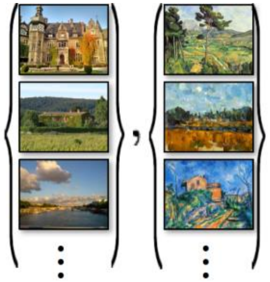

Generate

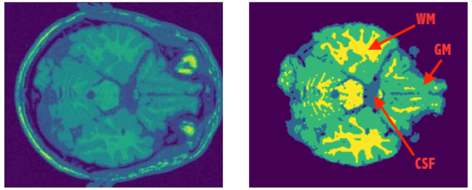

Segment, Detect

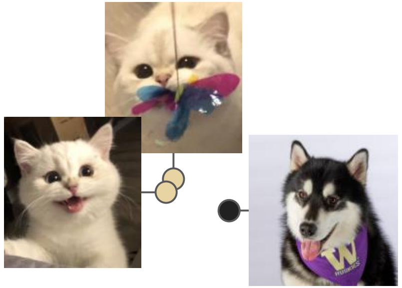

Contrastive

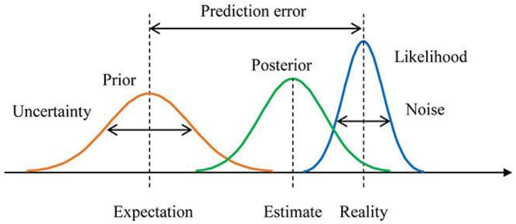

f(Confidence Esti.)

Medical Imaging

Ducts: for Breast Cancer Diagnosis

- **Important** Region for Diagnosis
- **Interpretable** Intermediate Results
- Only **240** Images (Biopsies) in the Dataset
- **No** Duct **Label**

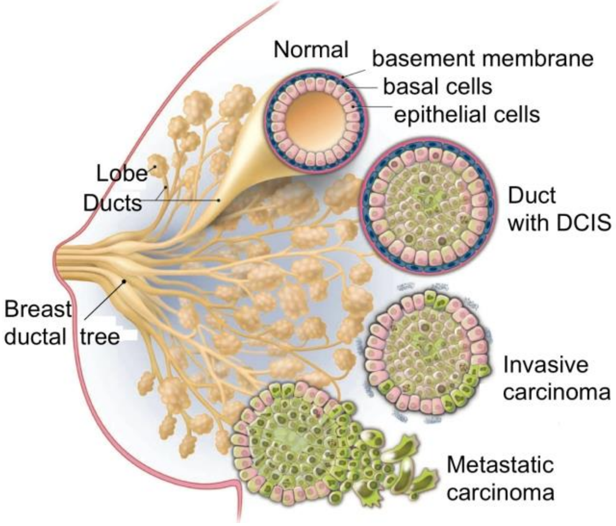

Classifying Breast Histopathology Images with a Ductal Instance-Oriented Pipeline
Li, B.; Mercan, E.; Mehta, S.; Knezevich, S.; Arnold, C.; Weaver, D.; Elmore, J.; Shapiro, L.
In 2020 25th International Conference on Pattern Recognition. IEEE.

Ductal Instance Identification

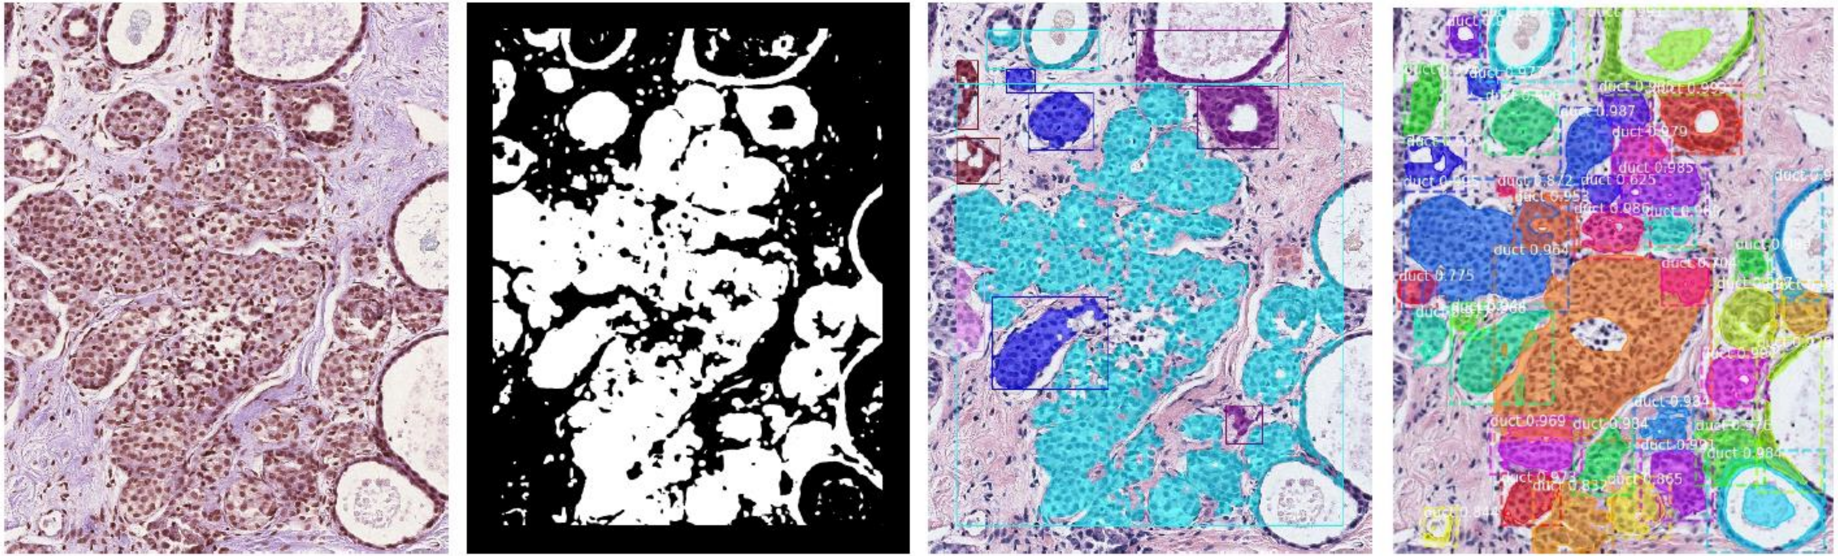

(a) ROI image (b) Binary (c) Previous (d) Ours

Weakly Supervised Annotation

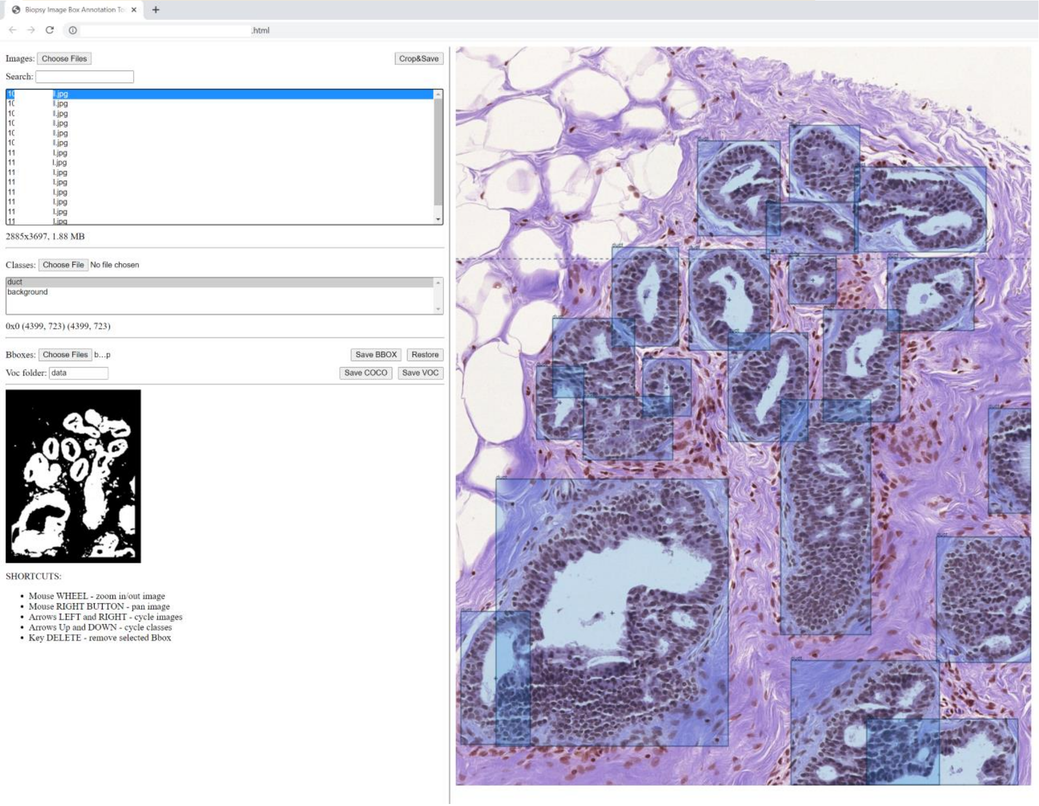

- Human-AI Collaboration
- Weakly Supervised Labeling
- Combines
  - Machine’s Semantic Prediction
  - Human Box Annotation
- Instance Segmentation Labels
- Silver Standard

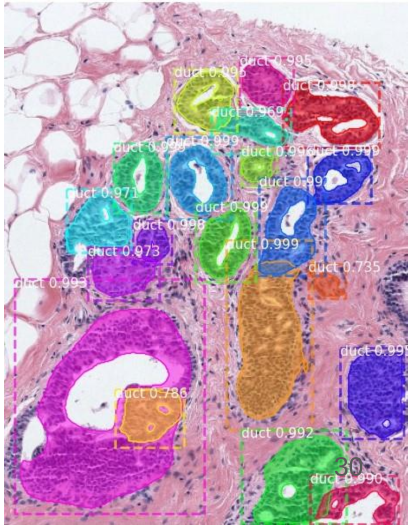

Diagnostic Performance

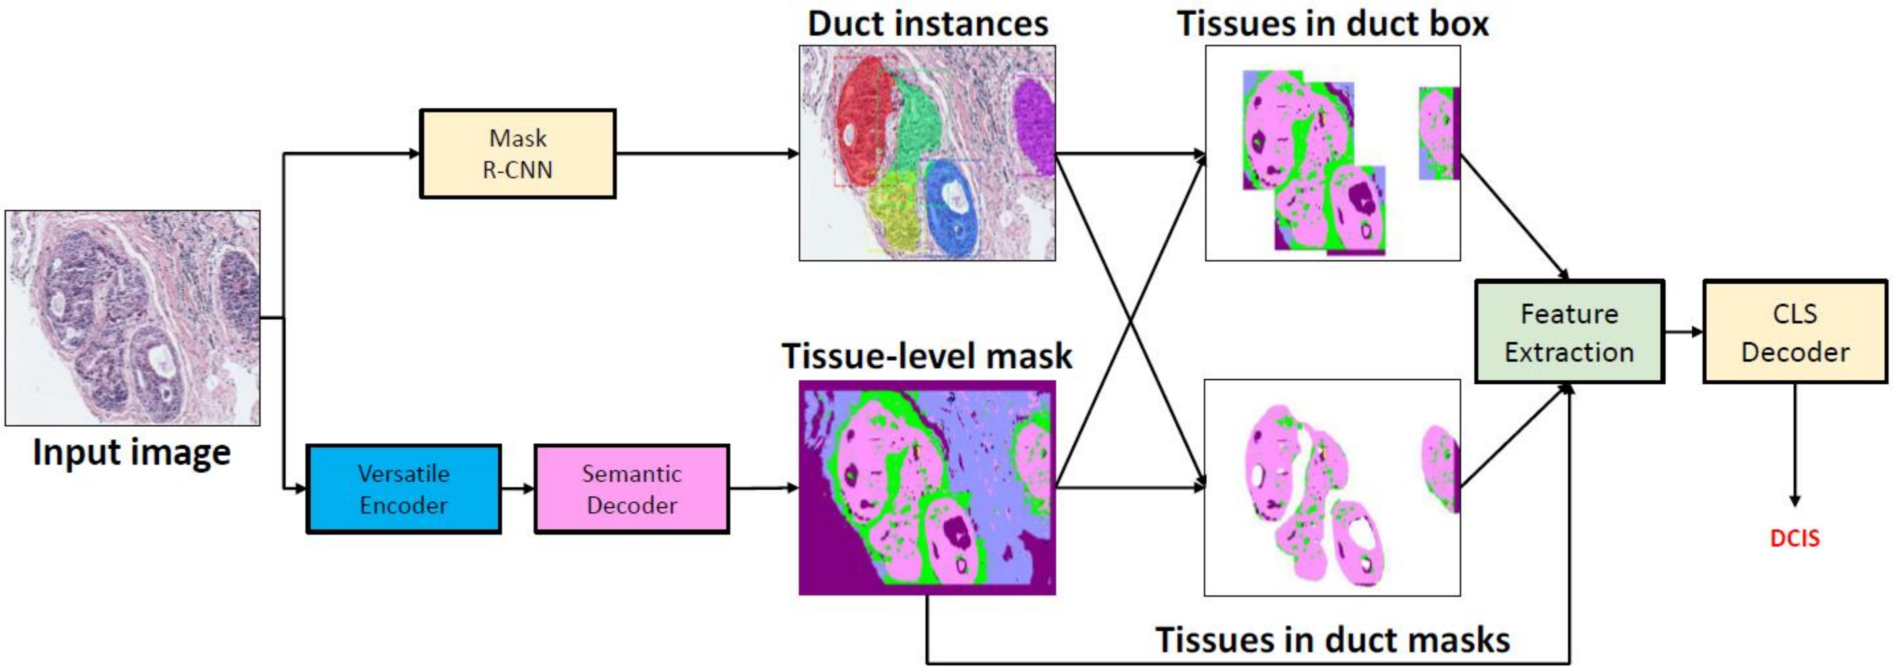

Instance Segmentation Results

Good

Imperfect

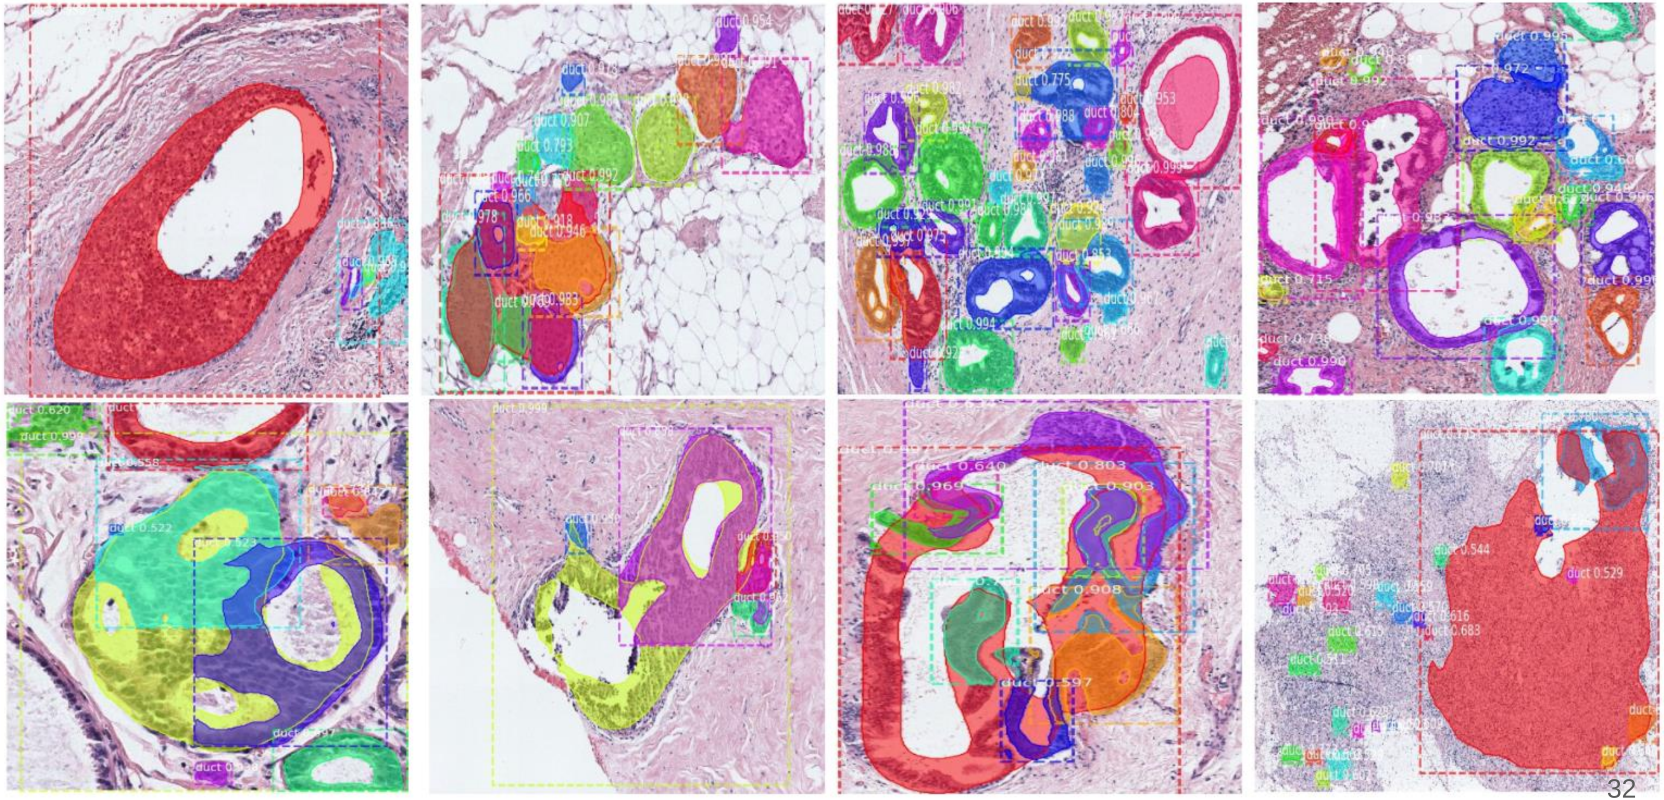

Diagnostic Results

Task	Features	Sensitivity	Specificity	Accuracy	F_1
Invasive vs Non-invasive	*Pathologists*	*0.84*	*0.99*	*0.98*	*0.86*
	Superpixel Features	**0.70**	0.95	0.94	0.62
	Structure Features	0.49	0.96	0.91	0.51
	Ours	0.62	**0.98**	**0.95**	**0.73**
Atypia and DCIS vs Benign	*Pathologists*	*0.72*	*0.62*	*0.81*	*0.51*
	Superpixel Features	0.79	0.41	0.70	0.46
	Structure Features	**0.85**	0.45	0.70	0.50
	Ours	**0.85**	**0.63**	**0.79**	**0.59**
DCIS vs Atypia	*Pathologists*	*0.70*	*0.82*	*0.80*	*0.76*
	Superpixel Features	0.88	0.78	0.83	0.86
	Structure Features	0.89	0.80	0.85	0.87
	Ours	**0.91**	**0.89**	**0.90**	**0.92**

Method	Accuracy
Pathologists	0.70
MIL with max-pooling	0.55
MIL with learned fusion	0.67
Semantic Learning	0.55
Y-Net	0.63
Ours	**0.70** $\pm$ 0.02

Ablation Experiments

Malignant Epithelium (BE) and Necrosis (NC) co-occurs

Benign Epithelium (BE) near Boundary (BD)

Method	Accuracy
Tissue in ROI	0.67
Tissue in Duct box	0.66
Tissue in Duct mask	0.69
Tissue in Duct mask + ROI	0.69
Tissue in Duct box + ROI	0.67
Tissue in Duct box + mask	0.69
Tissue (All)	**0.70**

Rank	DIOP (ours)	Tissue-level model
1	BD & BE in duct mask	ME & NC in ROI
2	ME & NC in duct mask	BG & NC in ROI
3	BD & NC in duct mask	SC freq in ROI
4	BE & NS in bounding box	BE freq in ROI
5	BG & NC in duct mask	BE & SC in ROI
6	BE & SC in ROI	ME & NS in ROI
7	ME & SC in bounding box	BE & NS in ROI
8	NC freq in bounding box	NC freq in ROI
9	BE & SC in bounding box	NS & NC in ROI
10	DS freq in duct mask	SC & NC in ROI

Necrosis (NC) near Boundary (BD)

Behavior Analysis

Data Collection

- From iPad Air 2 or iPad 5th Generation
- 88 participants
  - 49 children with ASD
  - 39 typically developing (TD) peers

Disclaimer: During our experiments, informed consent was obtained from parents of all children, and all study procedures were designed in accordance with the World Medical Association Declaration of Helsinki - Ethical Principles for Medical Research Involving Human Subjects as well as in compliance with HIPAA to preserve privacy. The Institutional Review Boards approved these studies of Yale University, Seattle Children's Research Institute, and the University of Washington.

Multi-Task Learning in Affective Computing

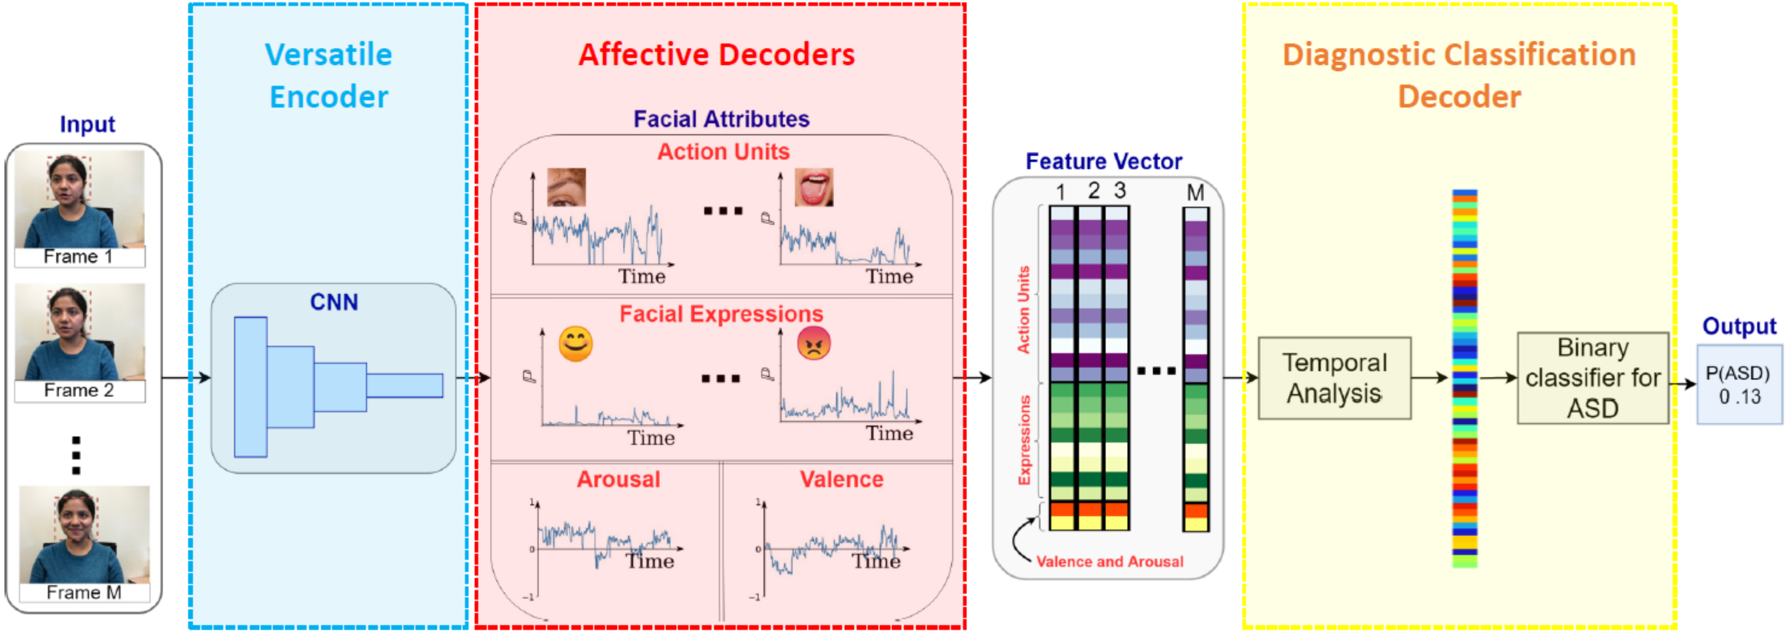

A Facial Affect Analysis System for Autism Spectrum Disorder
Li, B.; Mehta, S.; Aneja, D.; Foster, C.; Ventola, P.; Shic, F.; Shapiro, L.
In Proceedings of the IEEE International Conference on Image Processing (ICIP 2019)

Results

CNN Unit	# Params	FLOPs	Expression (F1)	AU (mF1Acc)	Valence (CC)	Arousal (CC)
Single-task						
Bottleneck	25.9 M	3.4 B	0.56	**0.78**	0.63	0.54
MobileNet	24.8 M	3.1 B	0.57	0.77	0.64	0.52
EESP	9.7 M	1.2 B	0.57	0.76	0.64	0.52
Multi-task						
Bottleneck	6.5 M	0.85 B	**0.58**	0.75	0.68	0.61
MobileNet	6.2 M	0.78 B	**0.58**	0.75	0.68	**0.62**
EESP	**2.4 M**	**0.29 B**	**0.58**	0.75	**0.69**	0.61
Literature						
SOTA	-	-	0.58	-	0.66	0.54
Human Performance	-	-	0.61	*	0.82	0.57

- Benitez-Quiroz, Carlos Fabian, Yan Wang, and Aleix M. Martinez. "Recognition of Action Units in the Wild with Deep Nets and a New Global-Local Loss." *ICCV*. 2017.
- Benitez-Quiroz, C. Fabian, et al. "EmotioNet Challenge: Recognition of facial expressions of emotion in the wild." *arXiv preprint arXiv:1703.01210* (2017).
- Mollahosseini, Ali, Behzad Hasani, and Mohammad H. Mahoor. "AffectNet: A Database for Facial Expression, Valence, and Arousal Computing in the Wild." *IEEE Transactions on Affective Computing* (2017).

Gaze Estimation (Face to Gaze)

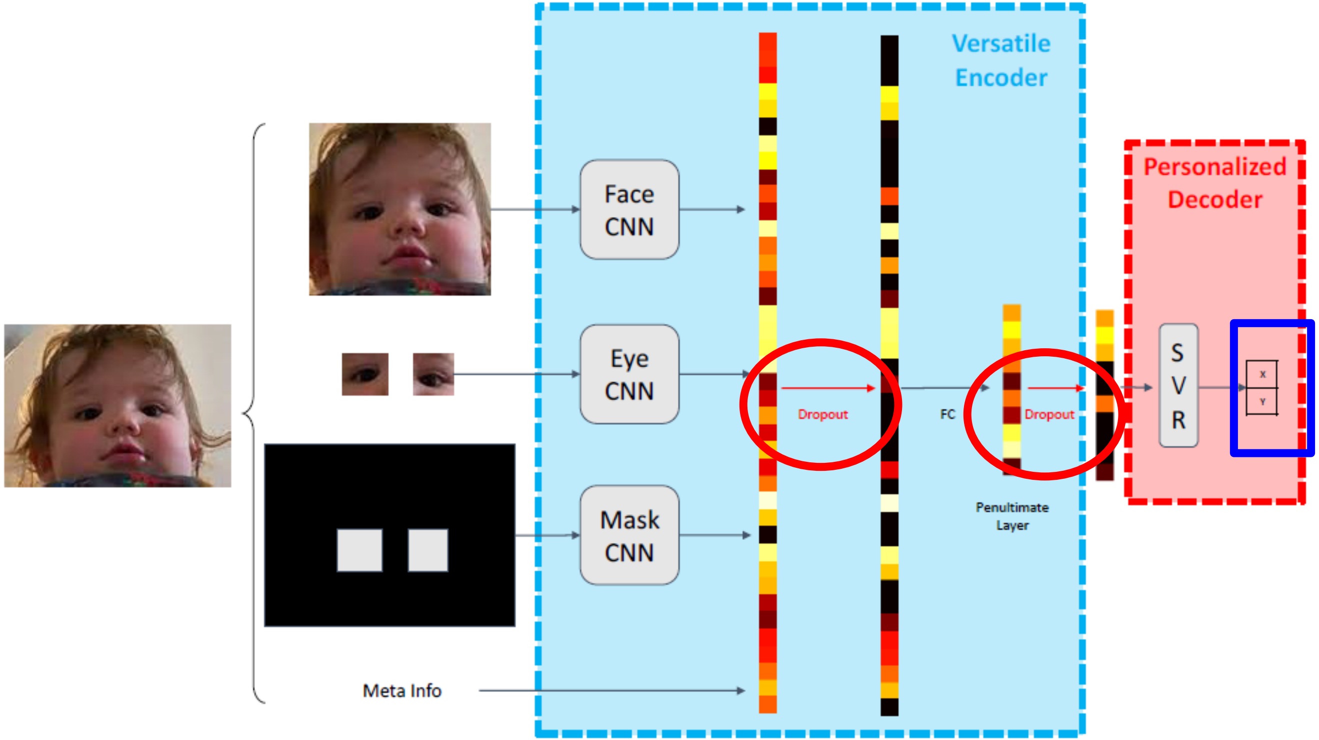

Severity of Data Shift (Post-Hoc)

Population	Data Source	Euclidean Error	L/R Accuracy	U/D Accuracy
Adults	GazeCapture	3.53	0.86	0.82
TD Children	Lab	4.65	0.82	0.74
Children with ASD	Lab	4.86	0.81	0.68
Children with ASD	Home	4.96	0.74	0.65

Adaptation (Calibration) Performance

Calibrate On	Euclidean Error	L/R Accuracy	U/D Accuracy
None	4.80	0.81	0.71
5-point	4.59	0.81	0.74
Smooth Pursuit	4.47	0.82	0.69
Both	4.32	0.82	0.72

Scanpath Analysis

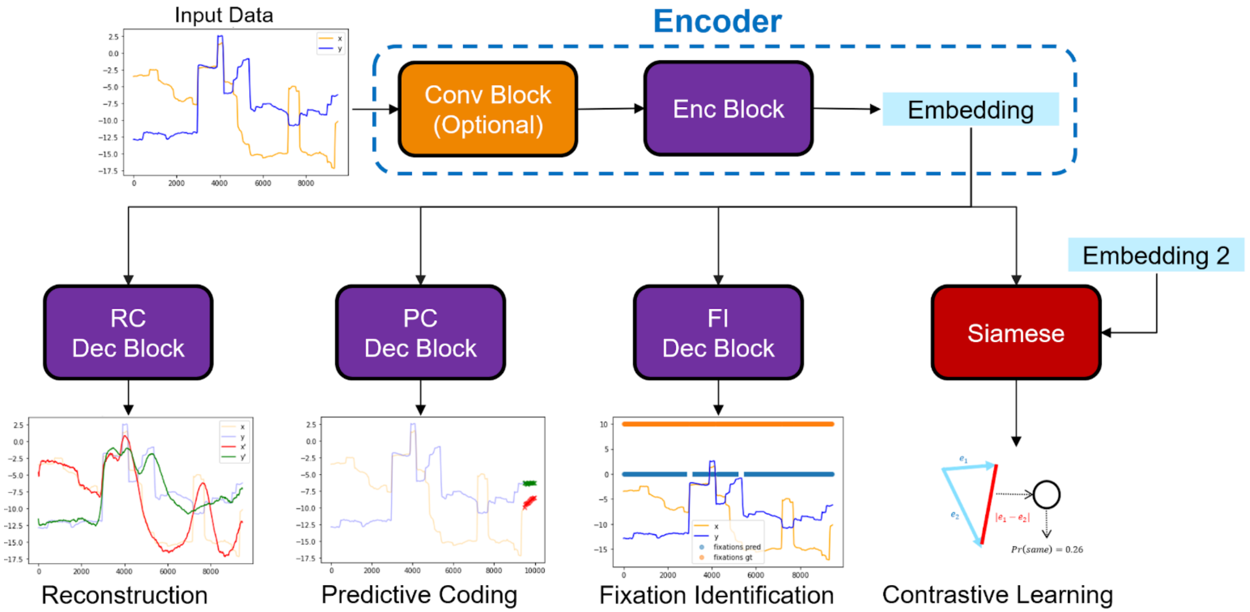

Learning Oculomotor Behaviors from Scanpath Data. **Li, B.**; Nuechterlein, N.; Barney, E.; Foster, C.; Kim, M.; Mahony, M.; Atyabi, A.; Feng, L.; Wang, Q.; Ventola, P.; Shapiro, L.; Shic, F. In 2021 *ACM International Conference In Multi-modal Interaction (ICMI).*

Application Group Classification

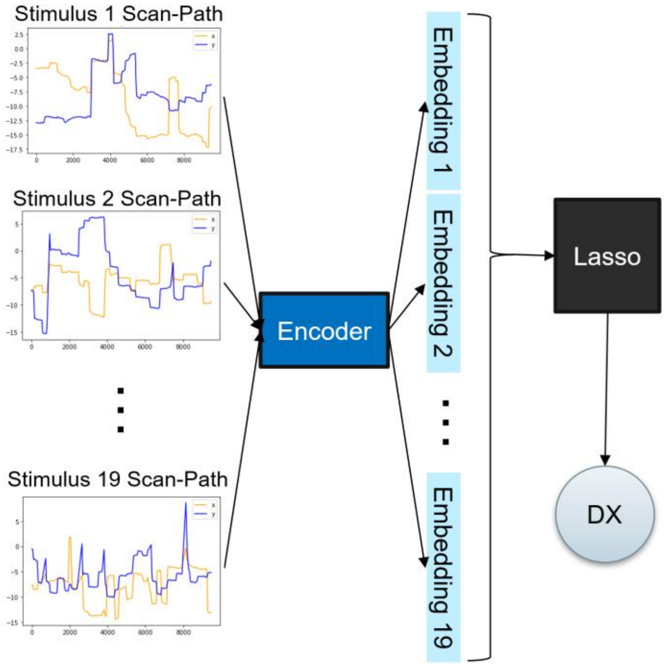

Method	Accuracy	AUC	F-1
Expert Features from	0.74	0.68	0.82
SGIN (with more data)	0.78	**0.83**	0.83
Ours	**0.80**	**0.83**	**0.88**

Results for Autism Classification

Query Optimization in DBMS

Cardinality Estimation

```
SELECT p.Name AS ProductName,
NonDiscountSales = (OrderQty * UnitPrice),
Discounts = ((OrderQty * UnitPrice) *
UnitPriceDiscount)
FROM Production.Product AS p
JOIN Sales.SalesOrderDetail AS sod
ON p.ProductID = sod.ProductID
ORDER BY ProductName DESC;
```

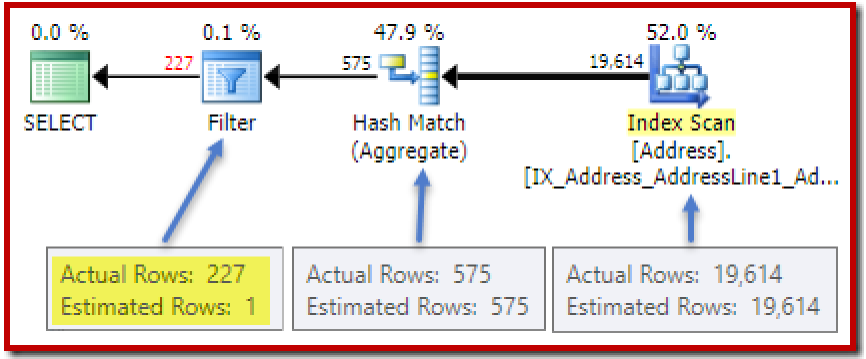

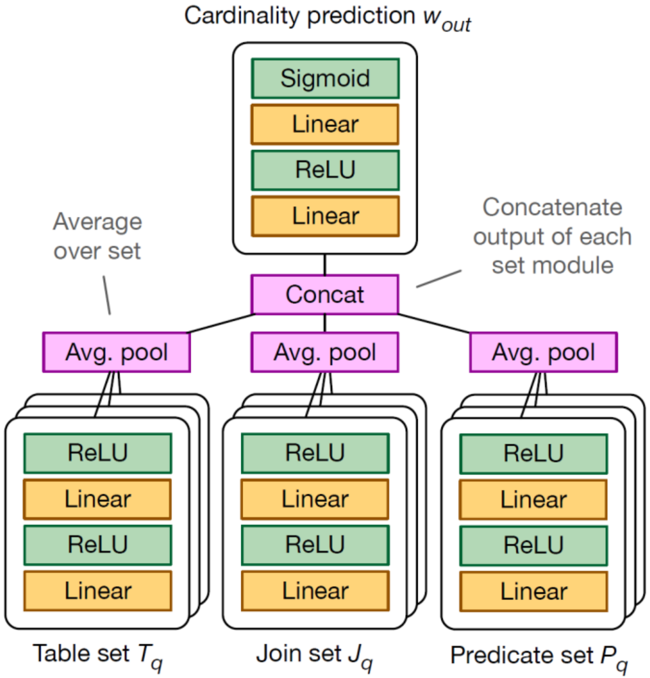

- Cardinality Estimation: Is Machine Learning a Silver Bullet? Li, B.; Lu, Y.; Wang, C.; Kandula, S.. The 3rd International Workshop on Applied AI for Database Systems and Applications (AIDB).
- Q-error Bounds of Random Uniform Sampling for Cardinality Estimation. Li, B.; Lu, Y.; Wang, C.; Kandula, S..
- Warper: Efficiently Adapting Learned Cardinality Estimators to Data and Workload Drifts. Li, B.; Lu, Y.; Wang, C.; Kandula, S..

Data Shift After Deployment

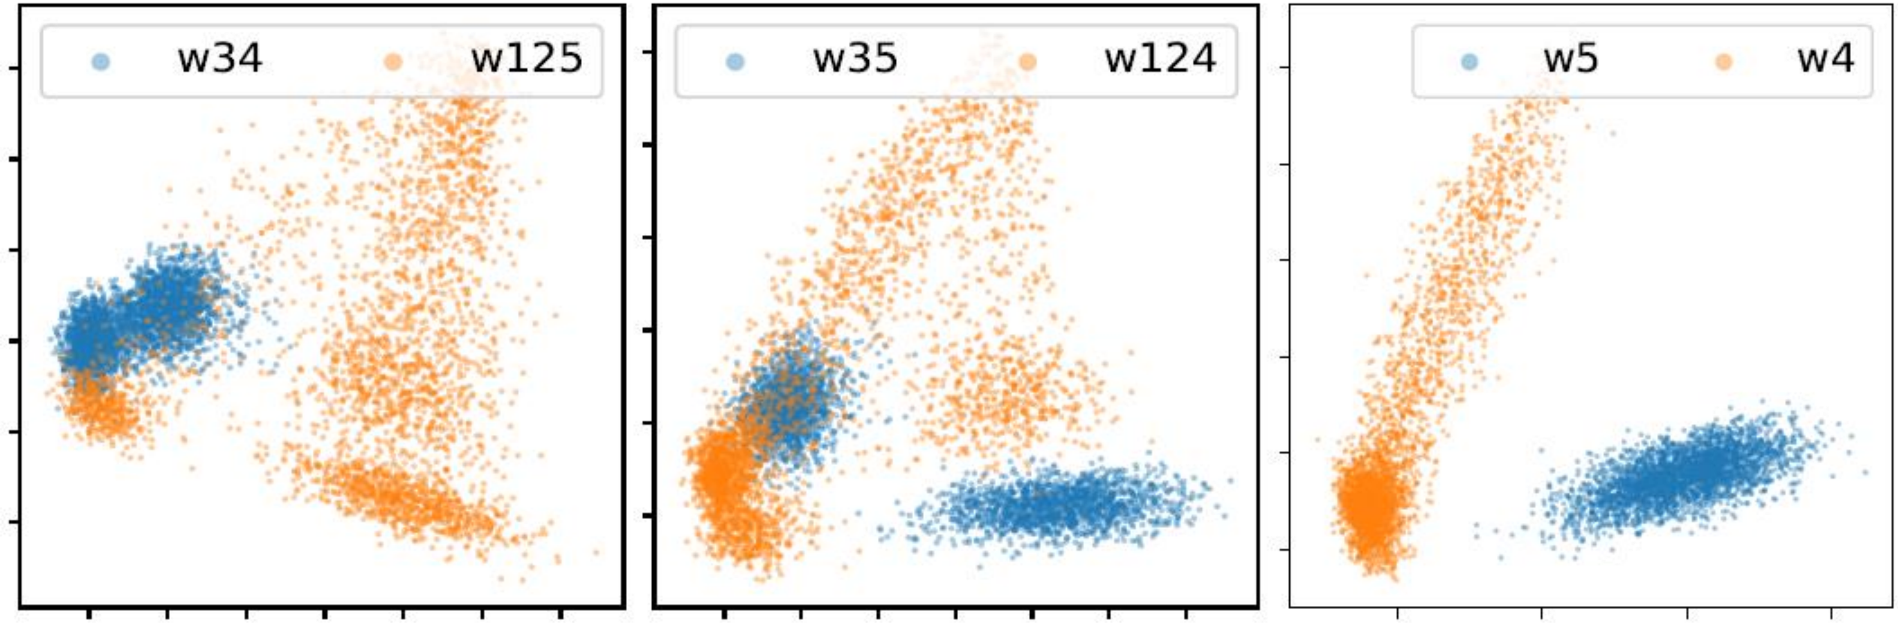

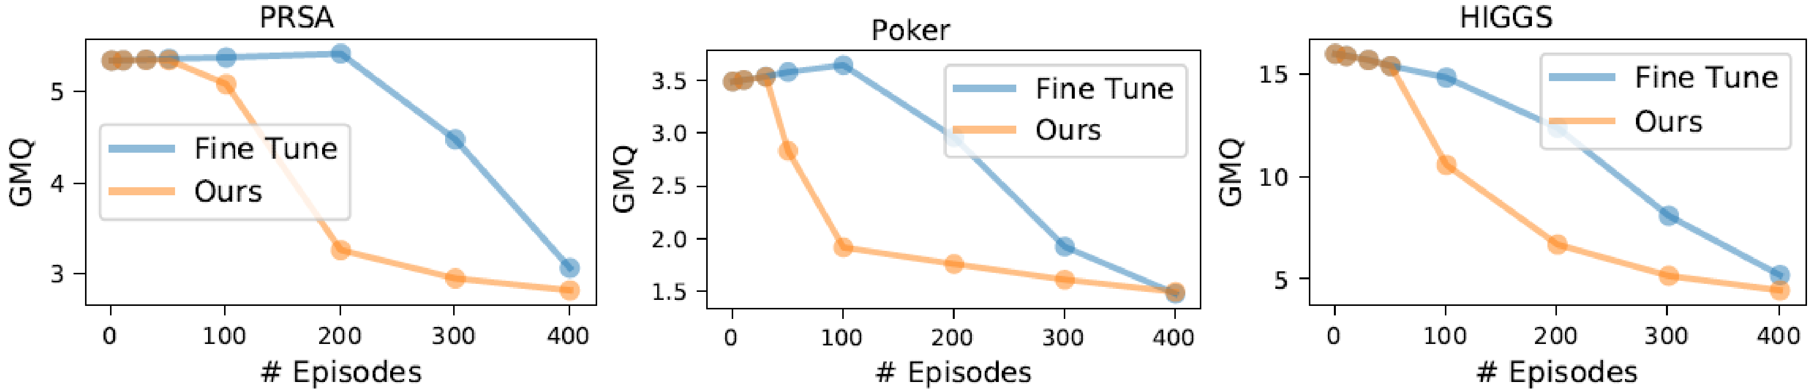
PRSA
Poker
HIGGS
Fine Tune
Ours
GMQ
Episodes

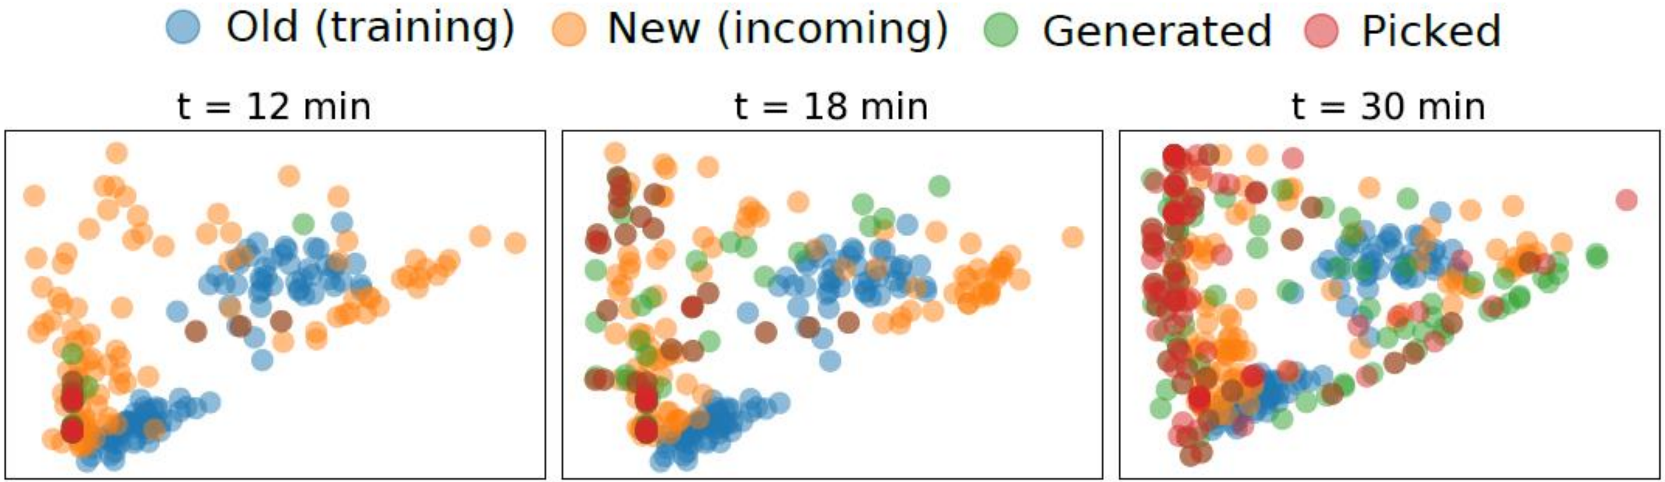
Old (training)
New (incoming)
Generated
Picked
t = 12 min
t = 18 min
t = 30 min

Discussion

Future Work

- Apply and Test on More Modalities (NLP, DNA, ...)
- More Theoretical Analyses
- Multi-Encoder
- Improve User Interface for AutoML
- Federated Learning, Decentralized Learning
- Diversity, Equality, Fairness, and Bias

Next …

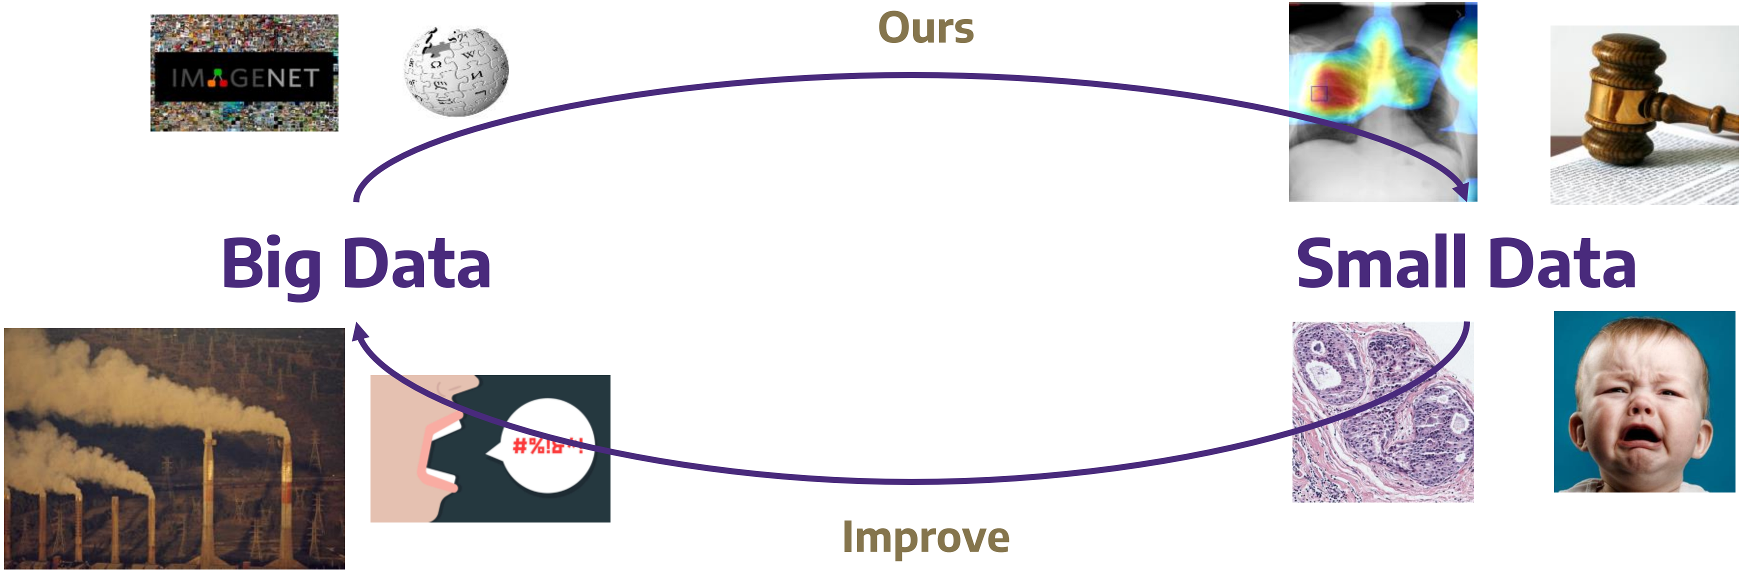

Thanks!