Challenges in building AI systems for Smart Health

http://mialab.net

SAMSON CHEUNG, ECE DEPARTMENT
UNIVERSITY OF KENTUCKY

Research Interests

- Professor at UK since 2004
- Image Processing
  - Color, 3D, thermal images and video processing
- Security & Privacy
  - Encrypted-domain signal processing
  - Differential Privacy
- Applied Deep Learning
  - Generative Models
  - Bayesian Modeling
  - Applications in Smart Health
- Technology based Autism Research
  - Assistive technologies
  - Autism screening

Outline

- AI Challenges on Smart Health
- Autism Risk Prediction
- Whole Slide Image Segmentation
- Data privacy in Machine Learning
- Conclusions

Challenges to apply AI in Health

Big BIG Data

Annotation is Expensive

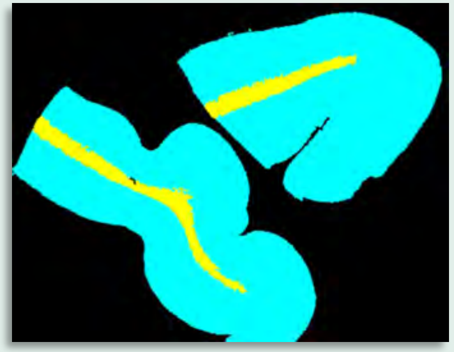

Overemphasis on Supervised Learning

Labeled Data → Supervised Learning

Unlabeled Data → Active Learning

Labeled and Unlabeled Data → Semi-Supervised Learning

Data Privacy

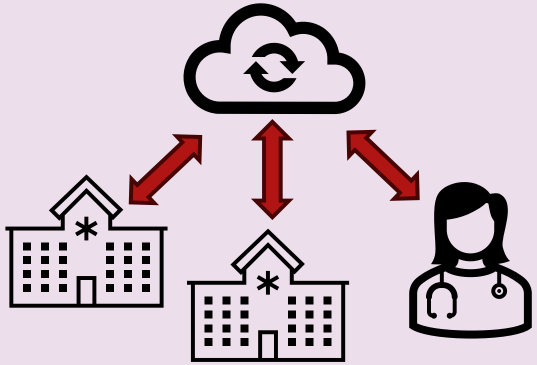

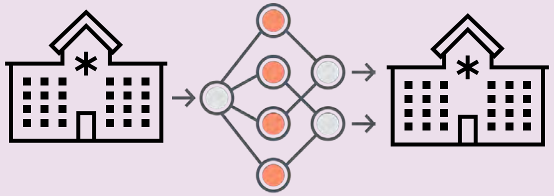

Autism Risk Prediction based on behavior markers

Autism Spectrum Disorder (ASD)

- What is it?
  - Neuro-developmental disorder
  - Significant social, communication and behavioral challenges
- What is the societal impact?
  - 1 in 36 children in the US diagnosed (CDC, 2023)
  - Total lifetime cost = $11.5 trillion dollars by 2029 (Cakir et al., 2020)
  - Early intervention is important for optimal outcome
  - Average diagnosis at 60 months (NSCH 2019)

Repetitive behavior

Poor eye contact

Self-injurious behaviors

ASD Risk from Dyadic Behaviors

- UC Davis Infant Sibling Study (2003 – 2023)
- Interaction between an adult and a child
- 547 subjects: 6, 12, 18, 24 and 36 months
- Concurrent diagnosis: 60 subjects are ASD
- Over 300,000 minutes of video
- Manually coded behavior labels: look-face, look-object, smile, vocalization

ASD Risk from Dyadic Behaviors

Video Frames → **Different CV+Audio features** → **Behavior Classifier** → **Behavior Statistics** → **ASD Classifier** → ASD?

look-face, look-object, smile, vocalization

Behavior Statistics (tabular data)
- Frequency of behaviors
- Duration of behaviors
- Gender
- Age

ASD Classifier
- 3-layer MLP neural network
- SMOTE & Tomek Links for class balancing

Expert Labels (blue) vs Model Predictions (red) for Look Face Classifier

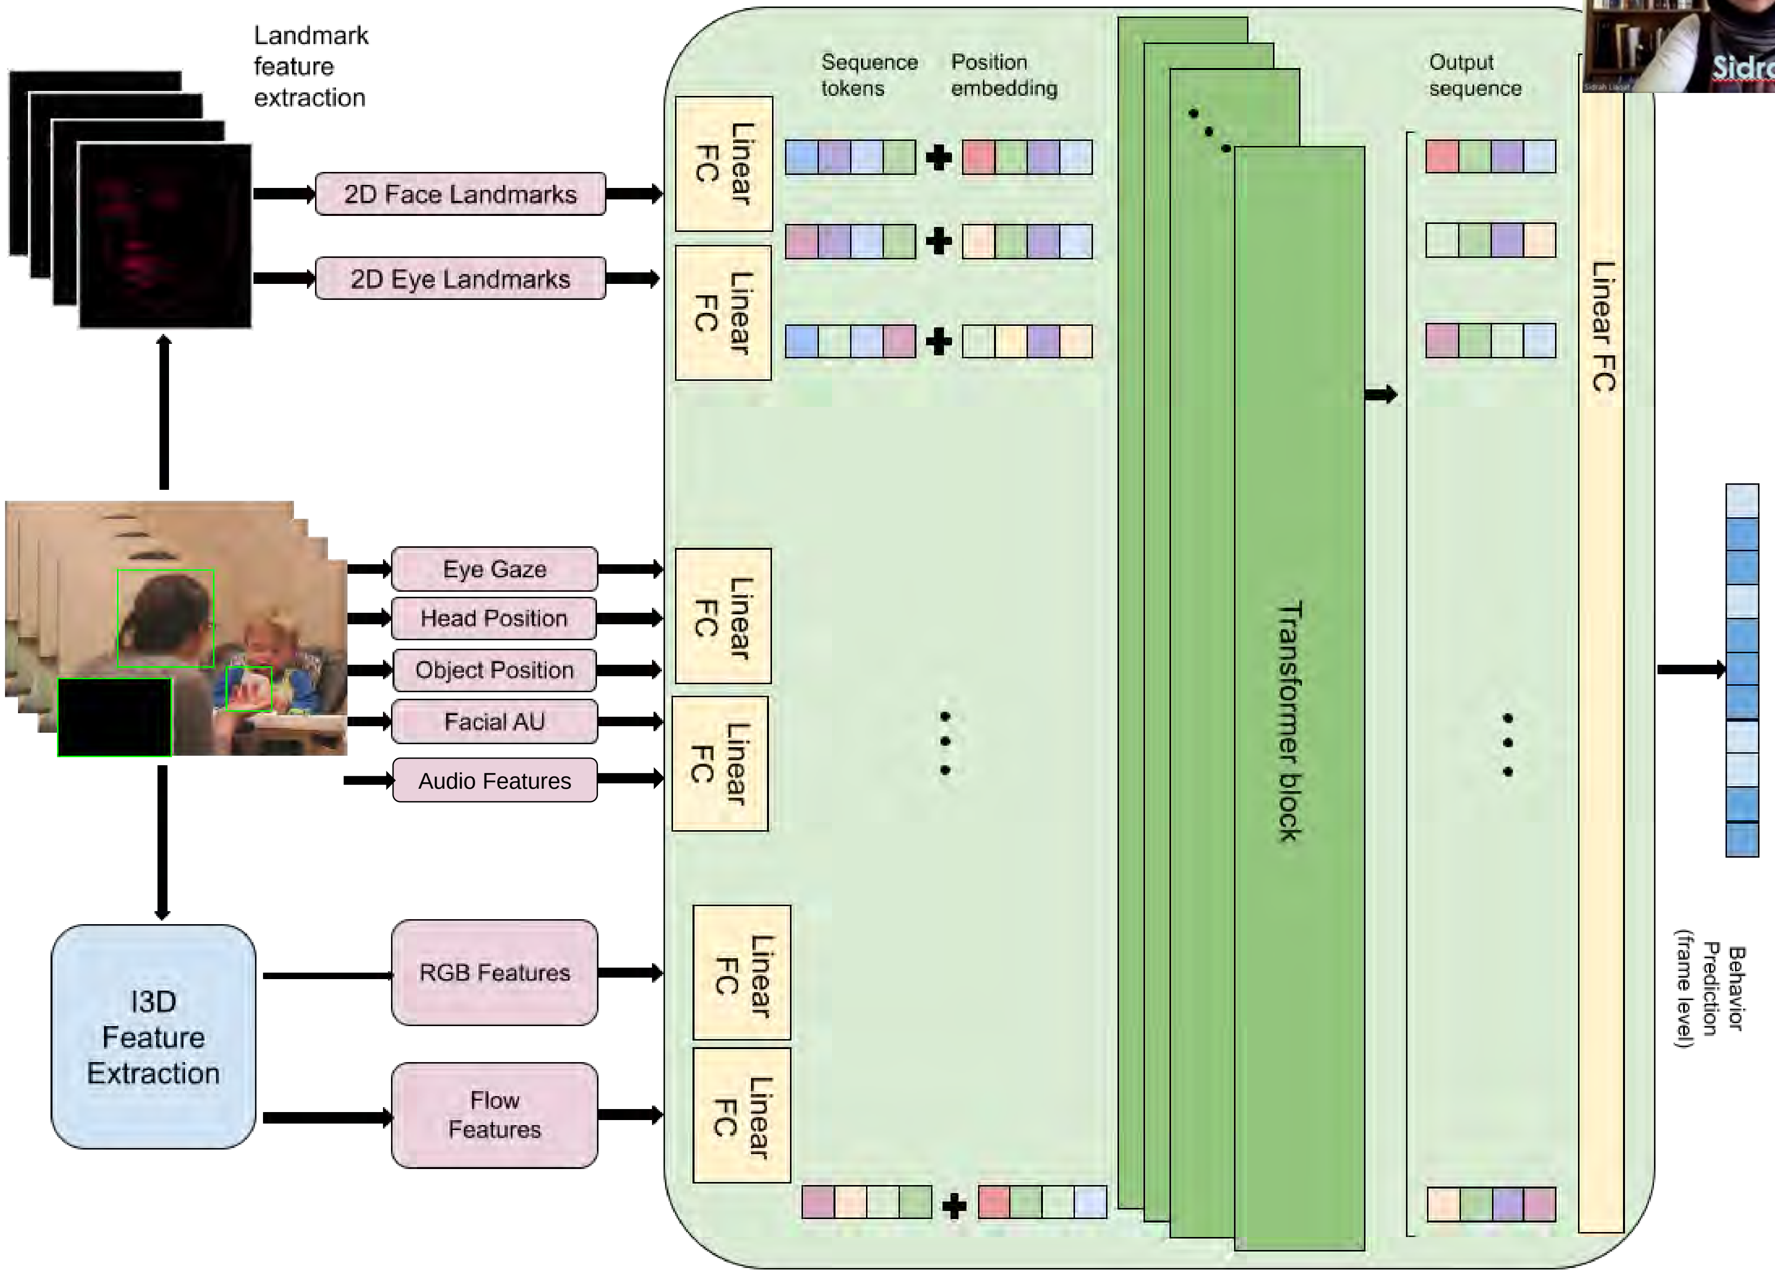

Multi-head self-attention based Transformer model

Detailed Architecture

- Expert Features: 2D facial and eye landmarks, facial action units, gaze direction, head and object locations
- Deep-learned features
  - I3D: mage and motion feature
  - Audio Frequency Mel Spectrum
- Short-time Transformer Architecture

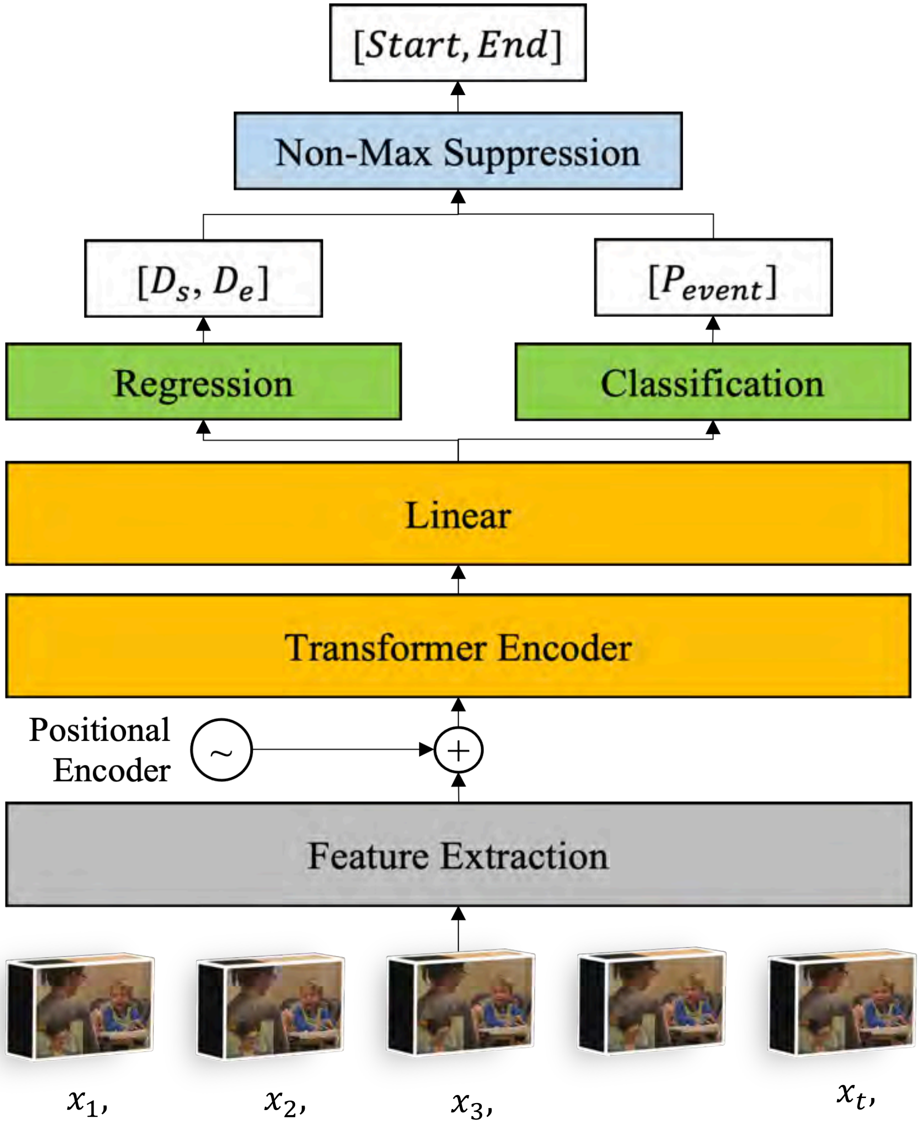

Temporal Action Localization

Alternative head: Frame prediction to segment detection

- **Frame-based Methods:** $\mathbf{Y} = \{C_t\}$
  - $\boldsymbol{C_t}$: Frame behavior prediction
- **Segment Detection:** $\mathbf{Y} = D_t^S, D_t^E, p(C_t)$
  - $\boldsymbol{D_t^S}$: Distance to Start
  - $\boldsymbol{D_t^E}$: Distance to End
  - $\boldsymbol{P_{event}}$: Action Probability
- **Non-Max Suppression:**
  - Suppress unlikely behavior segment proposals

Experiments - Behavior Detection

Behavior Classification

	Sensitivity	Specificity	Accuracy	AUCROC
Smile	0.54	0.93	0.86	0.73
Look Face	0.66	0.84	0.81	0.75
Look Object	0.79	0.64	0.76	0.72
Vocal	0.63	0.91	0.87	0.77

There is no overlap in subjects between testing and training datasets.

ASD Risk Prediction

	Sensitivity	Specificity	Accuracy	AUCROC
Hand coded behavior	0.76	0.86	0.85	0.81
ML behavior	0.76	0.73	0.73	0.74

Sidrah Liaqat, et al. 2024. End-to-end Multi-Modal Behavior Based Autism Spectrum Disorder Detection from Video. In preparation..

Segmentation of Whole Slide Brian Tissue Image

Whole Slide Imaging in Pathology

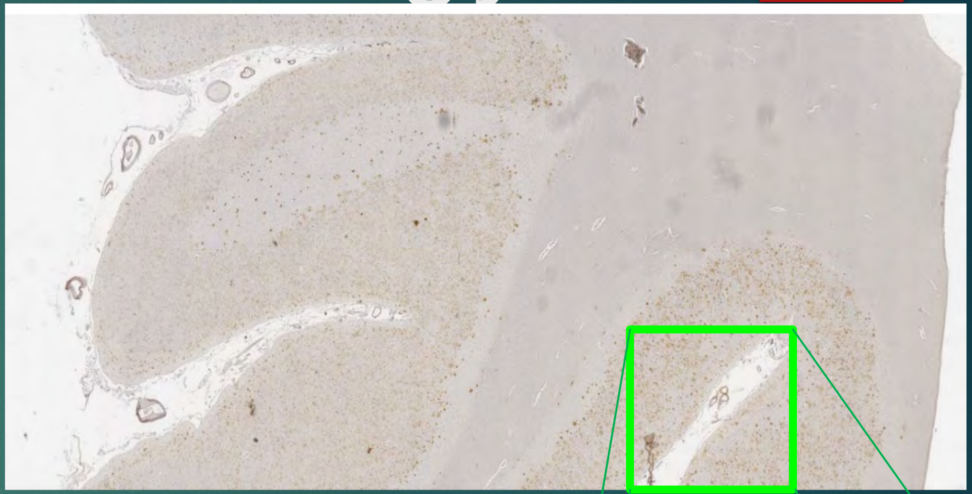

- Alzheimer's Disease (AD)
  - Most common cause of dementia
  - 6.9 million in US (1 in 9 age >65)
  - 24 million worldwide
- AD pathologies:
  - Amyloid beta (Aβ) plaques and cerebral amyloid angiopathy (CAA)
  - Found predominantly in Grey Matter (GM), less in White Matter (WM)
- Whole Slide Images (WSIs): brain tissue slides are stained & scanned with ultra-high resolution

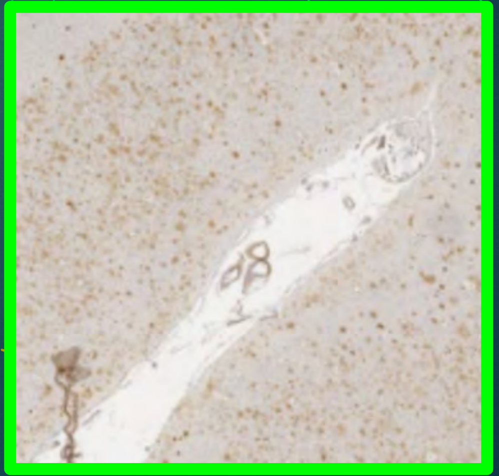

WSI Dataset

- Aβ stained WSIs
- AD: diagnosis of Alzheimer's disease
- NAD: No diagnosis of Alzheimer's disease
- 30 WSIs annotated by two trained personnel
- Resolution: nearly **60,000 × 50,000 (gigapixel)**

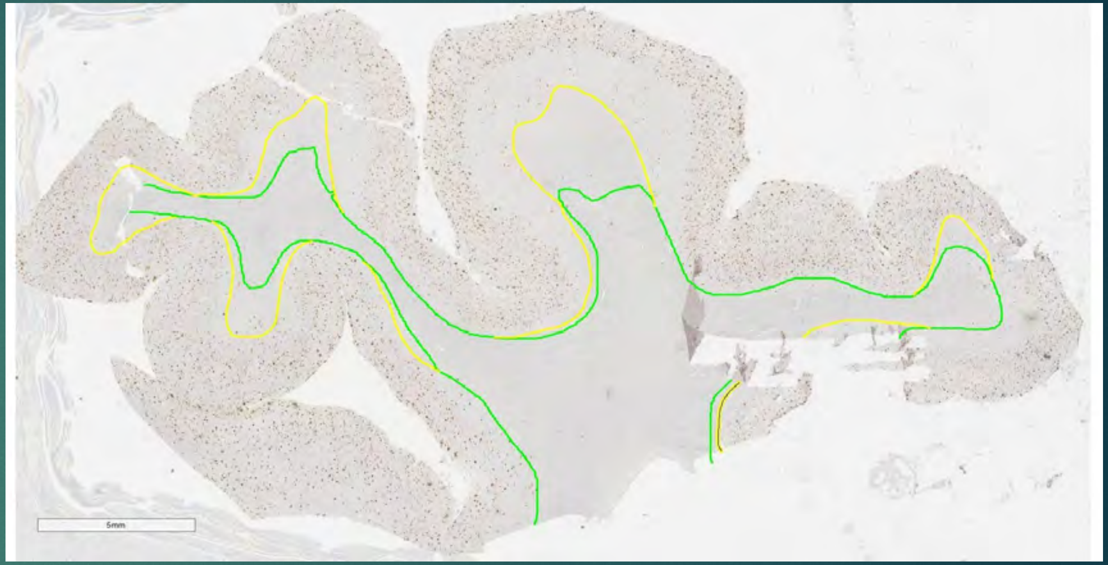

Data Split	AD	NAD
Training/Validation Set	12	8
Hold-out Test Set	6	4

GW/WM Basic Pipeline

Patch-wise Classification

Neural Conditional Random Field for post-processing

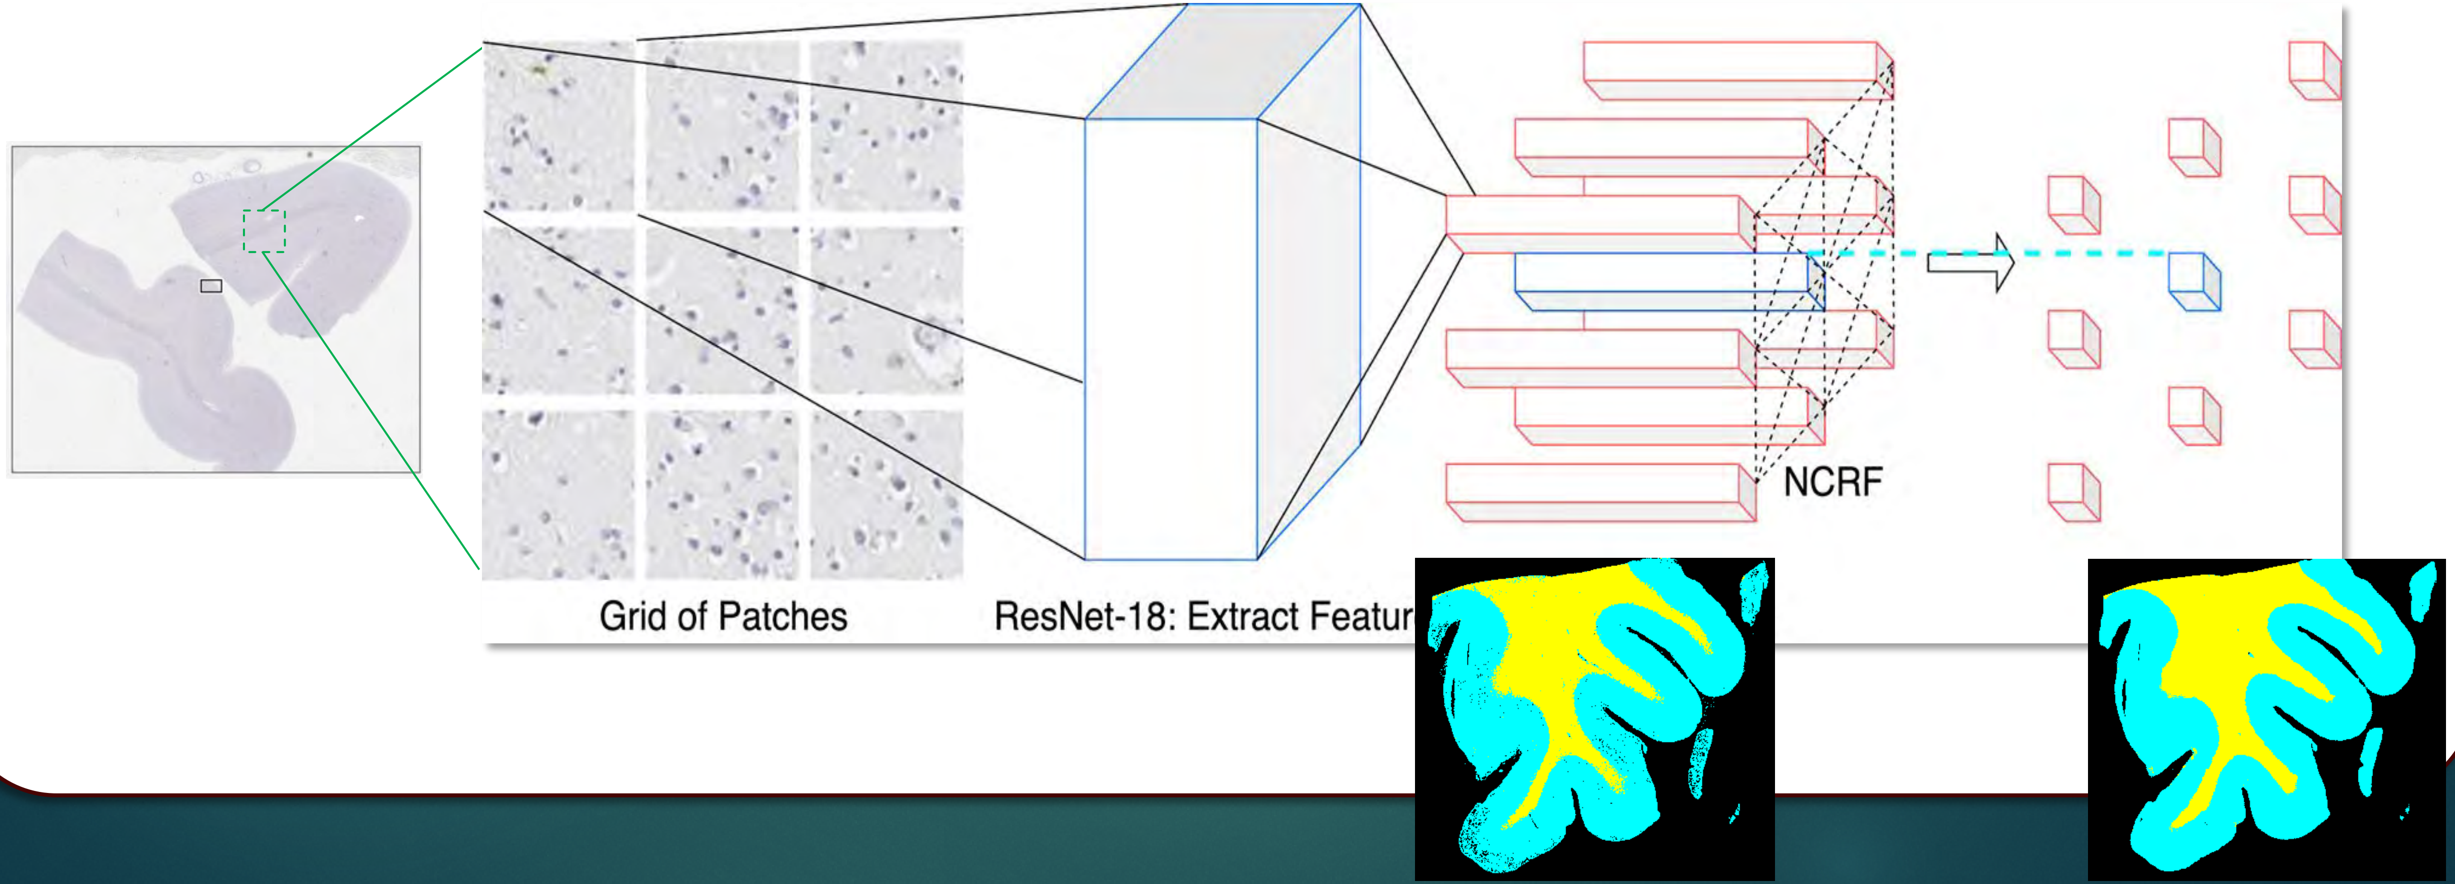

Prediction Masks

Ground Truth | FCN[2] | U-Net[3] | **ResNet-Patch** | **ResNet-NCRF**

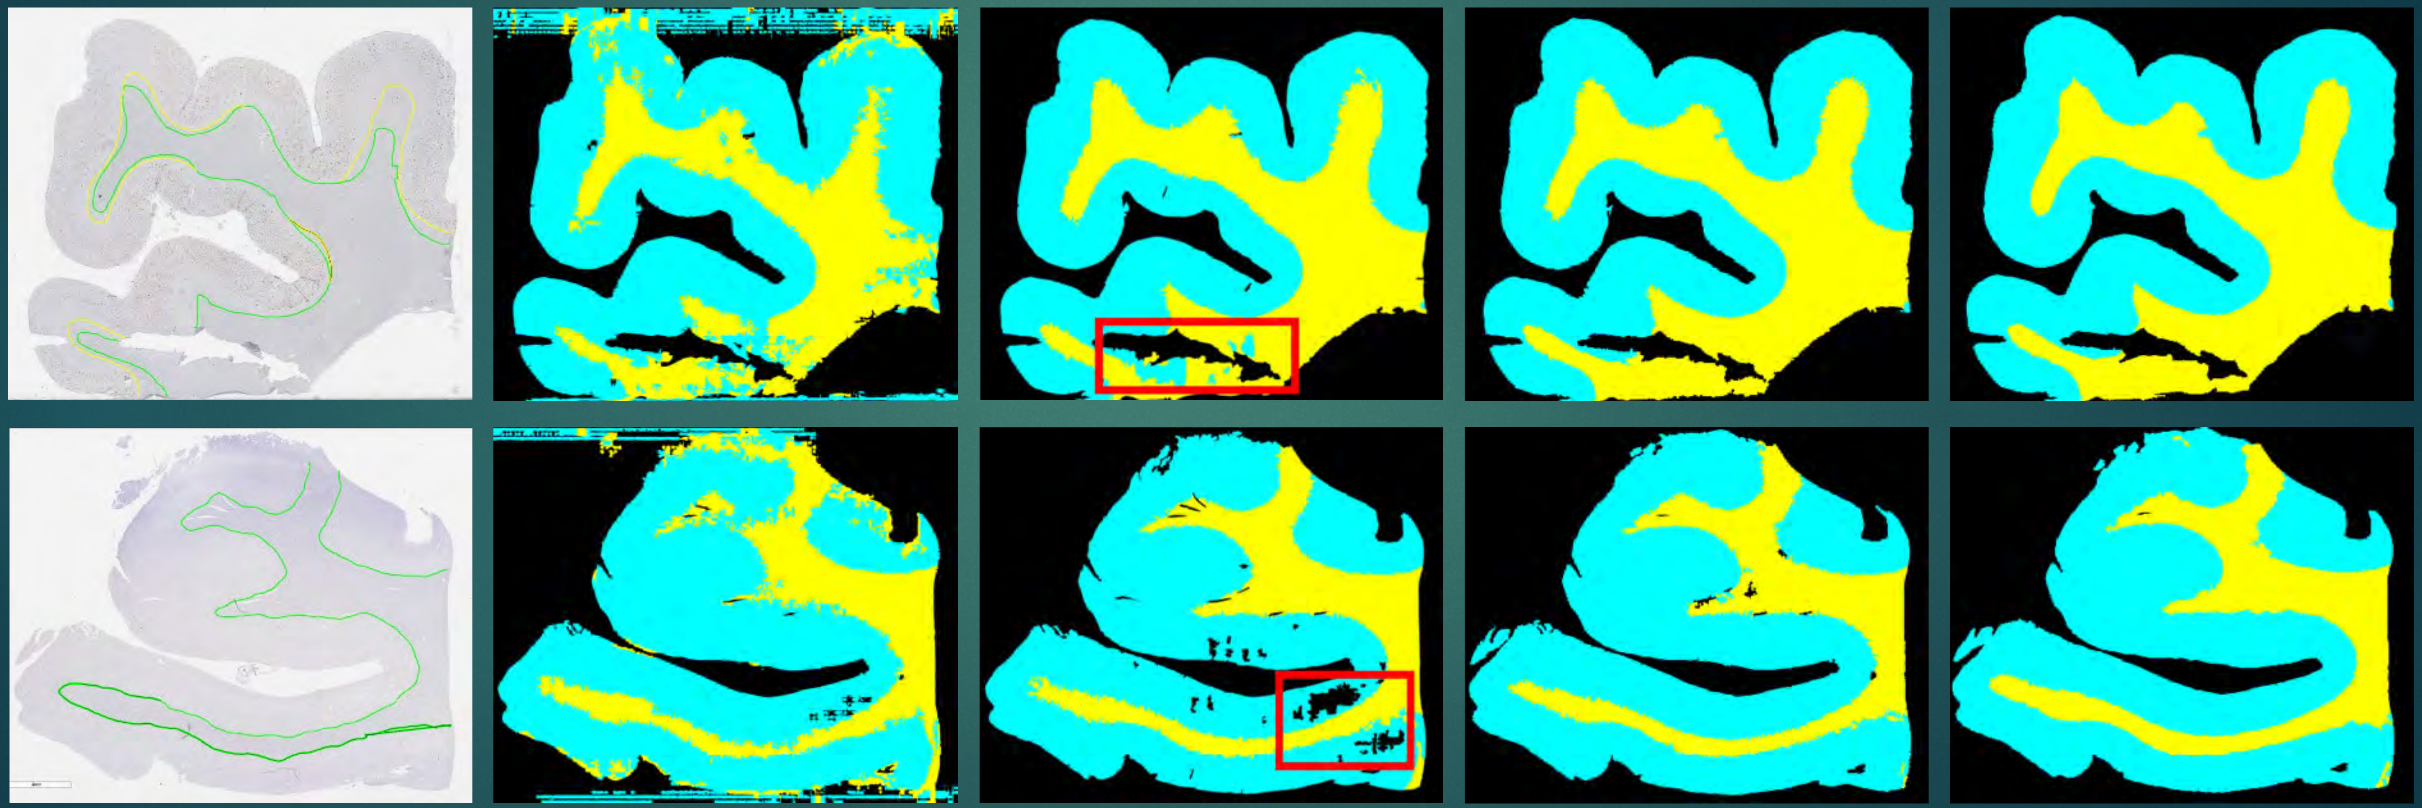

GM, WM, and background are indicated by cyan, yellow, and black, respectively.

Baseline Results

MEAN IOU ON 10 HOLD-OUT TEST WSIS (6 AD CASES AND 4 NAD CASES)

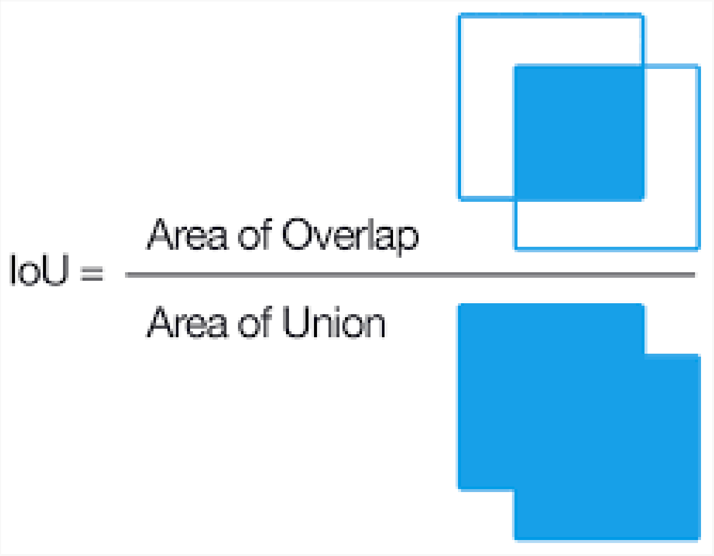

	FCN		U-Net		ResNet-Patch		ResNet
	Model	*Post*	*Model*	*Post*	*Model*	*Post*	-NCRF
Test Back	76.73 ± 8.57	84.05 ± 9.17	**97.25** ± 1.96	97.02 ± 2.15	97.06 ± **1.01**	97.10 ± 1.46	96.45 ± 1.95
Test GM	71.72 ± 8.93	77.39 ± 7.06	90.91 ± 4.90	91.52 ± 4.94	92.40 ± 2.83	**93.06** ± 2.71	**93.06** ± **2.20**
Test WM	49.09 ± 15.9	57.29 ± 14.3	79.12 ± 7.76	82.02 ± 7.98	81.36 ± 5.42	83.80 ± **5.28**	**84.27** ± 5.89
Test Mean	65.85 ± 9.33	72.91 ± 7.56	89.09 ± 3.71	90.19 ± 3.84	90.27 ± 2.13	**91.32** ± **1.94**	91.26 ± 1.99

Semi-supervised learning (SSL)

- Leverage unlabeled data to improve the performance when labeled data are limited
- FixMatch[1]
  - Consistency regularization
  - Pseudo-labeling (label assignment)
  - Combination of above
  - Achieve promising results when ***only use 40 labeled images*** in CIFAR-10

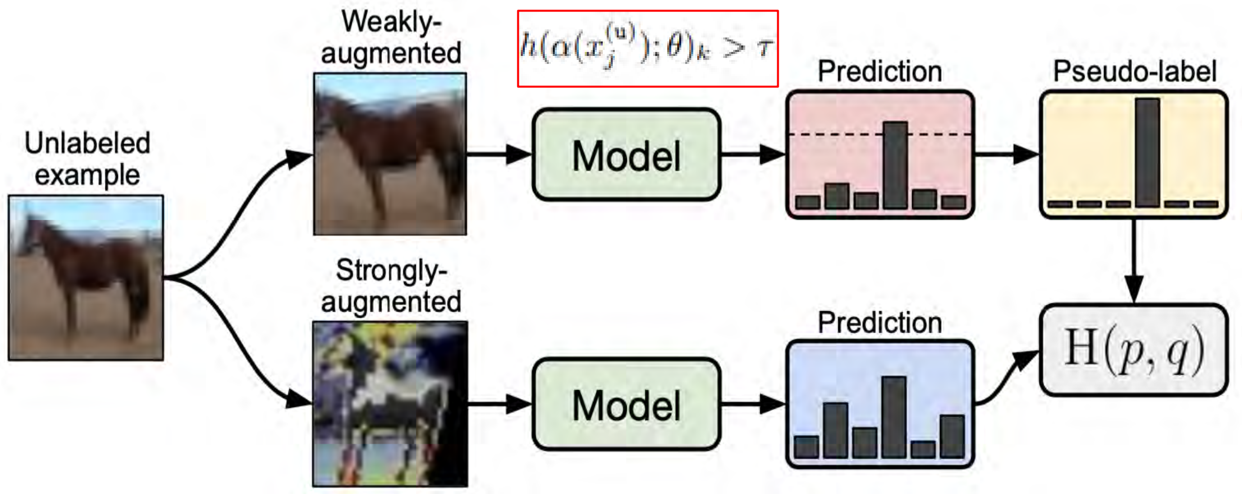

$$\min_{\theta\in\Theta} \sum_{i=1}^{m} \mathcal{L}(h(x_i^{(1)};\theta), y_i^{(1)}) + \sum_{j=1}^{n} \hat{\mathcal{L}}(h(\hat{\alpha}(x_j^{(u)});\theta), \hat{y}_j^{(u)}),$$

Supervised Loss **Consistency Loss**

[1] Sohn, et al. (2020). FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. Advances in Neural Information Processing Systems, 33.

Class imbalance on SSL

- SSL faces performance degradation when the unlabeled dataset is imbalanced

- Two problems:
  - *Confirmation bias* on pseudo labels – poor recall
  - *Mis-matched distributions* across the labeled, the unlabeled, and the test sets

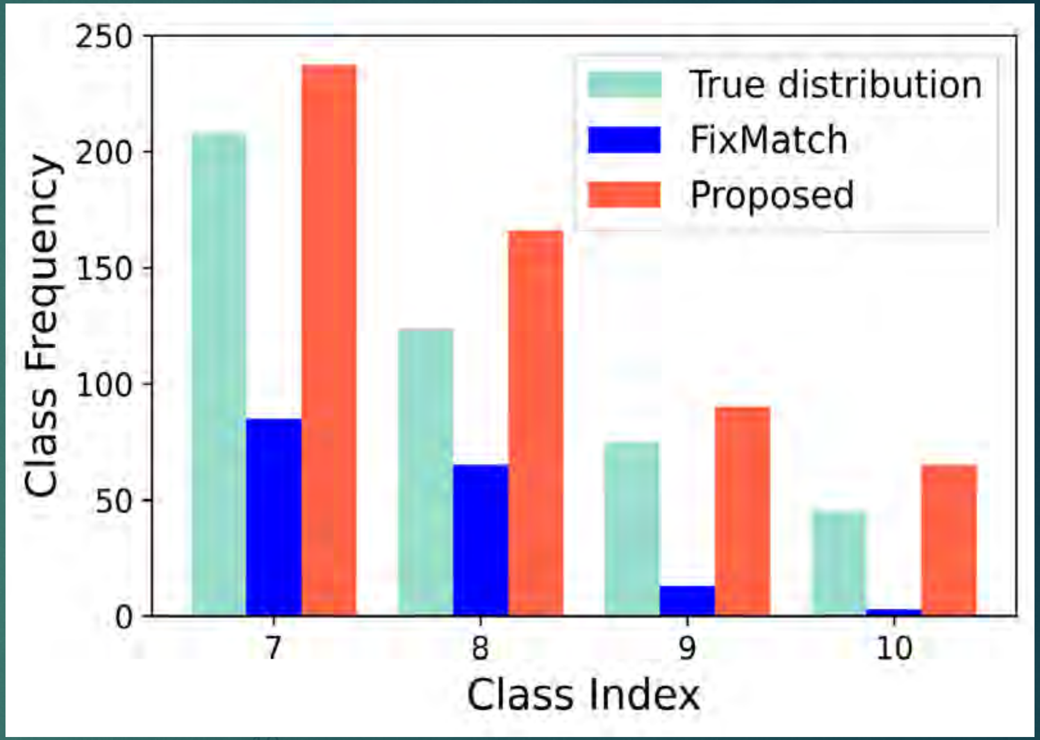

FixMatch underperforms on the minority classes (artifically skewed CIFAR10)

SaR: Self-adaptive Refinement on Pseudo-labels

Pseudo label refinement: $\hat{y}^{(u)}(k) = \delta(\boxed{w_k} \cdot h(\alpha(x^{(u)}); \theta)_k), k = 1, \cdots, C.$

Weighted consistency loss: $\mathcal{L}_{cw}(x; w, \theta) := \sum_{k=1}^{C} \boxed{w_k} \cdot p(x; \theta)_k \cdot \log(h(\texttt{pertub}(x); \theta)_k)$

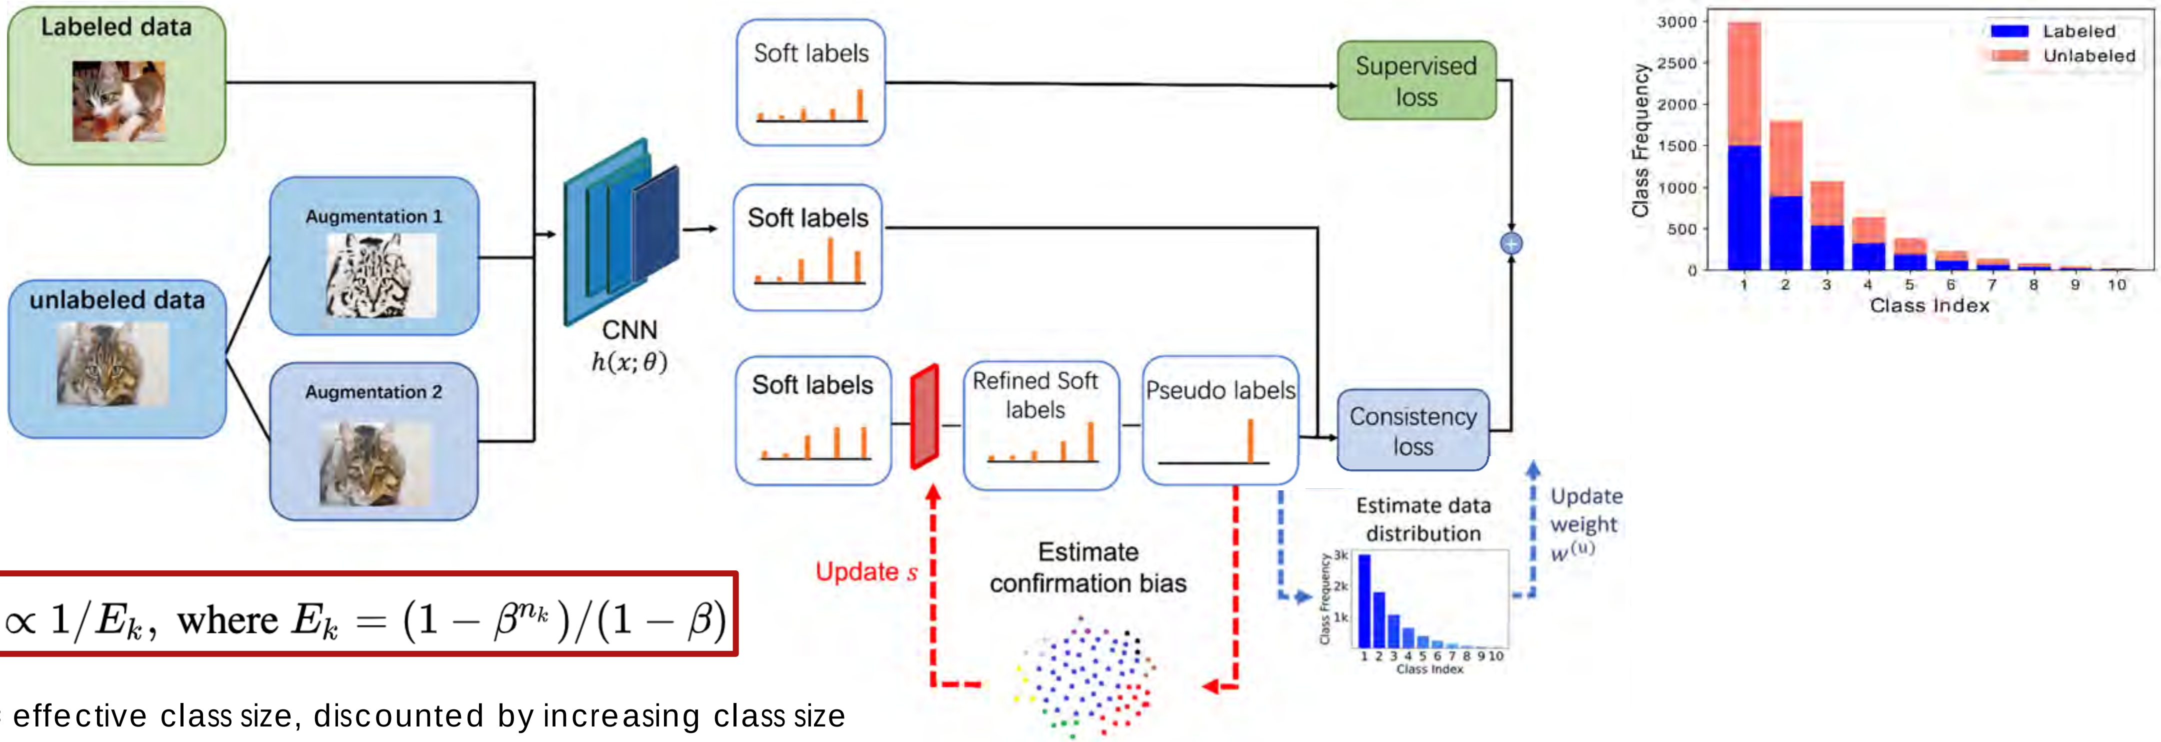

$$\boxed{w_k \propto 1/E_k,\ \text{where } E_k = (1-\beta^{n_k})/(1-\beta)}$$

E_k = effective class size, discounted by increasing class size

Results

I. U has a ***different*** distribution from L and the test set is balanced. (CIFAR-10)

Algorithm	$\gamma_u = 1$	$\gamma_u = 50$	$\gamma_u = 150$
ReMixMatch (Berthelot et al., 2020)	48.3±0.14 / 19.5±0.85	75.1±0.43 / 71.9±0.77	72.5±0.10 / 68.2±0.32
ReMixMatch* (Berthelot et al., 2020)	85.0±1.35 / 84.3±1.55	77.0±0.12 / 74.7±0.04	72.8±0.10 / 68.8±0.21
ReMixMatch* + DARP (Kim et al., 2020)	**89.7**±0.15 / **89.4**±0.17	77.4±0.22 / 75.0±0.25	73.2±0.11 / 69.2±0.31
ReMixMatch* + CReST (Wei et al., 2021)	45.9±1.27 / 20.1±1.99	70.2±0.45 / 65.8±0.71	65.4±0.34 / 62.9±0.15
ReMixMatch* + SAW	88.3±0.15 / 88.9±0.10	**80.3**±0.36 / **79.6**±0.40	**74.0**±0.94 / **72.4**±0.94
FixMatch (Sohn et al., 2020)	68.9±1.95 / 42.8±8.11	73.9±0.25 / 70.5±0.52	69.6±0.60 / 62.6±1.11
FixMatch + DARP (Kim et al., 2020)	**85.4**±0.55 / **85.0**±0.65	77.3±0.17 / 75.5±0.21	72.9±0.24 / 69.5±0.18
FixMatch + CReST (Wei et al., 2021)	60.2±1.34 / 35.9±2.50	65.8±0.78 / 67.1±0.84	60.1±1.44 / 51.4±1.68
FixMatch + SAW	83.9±0.44 / 83.3±0.47	**81.5**±2.25 / **80.9**±2.30	**76.8**±0.31 / **75.4**±0.37

- Measuring metric:
  - bACC (balanced accuracy)
  - GM (geometric mean)
- Imbalanced ratios (γ): for balanced set, it is set as 1.

II: U has a ***different*** distribution from L and the test set is ***imbalanced*** and of ***reversed*** distributions. (CIFAR-10)

Algorithm	$\gamma = 50$	$\gamma = 100$	$\gamma = 150$
ReMixMatch (Berthelot et al., 2020)	71.0±0.55 / 83.5±0.29	54.7±0.51 / 74.4±0.47	41.5±1.69 / 66.4±1.22
ReMixMatch + DARP (Kim et al., 2020)	66.9±0.75 / 80.5±0.46	49.7±1.55 / 70.5±0.90	35.8±1.81 / 60.9±2.42
ReMixMatch + CReST (Wei et al., 2021)	64.3±0.25 / 75.7±0.34	51.2±0.92 / 72.1±0.85	39.2±1.46 / 65.8±1.88
ReMixMatch + SAW	**86.3**±0.61 / **86.1**±0.64	**77.0**±0.59 / **76.0**±0.42	**71.5**±0.30 / **68.9**±0.26
FixMatch (Sohn et al., 2020)	70.5±0.26 / 82.2±0.31	51.0±1.65 / 71.5±1.24	38.5±1.15 / 63.4±0.31
FixMatch + DARP (Kim et al., 2020)	72.2±0.62 / 82.8±0.17	57.6±0.36 / 74.8±0.48	46.5±1.26 / 68.1±0.10
FixMatch + CReST (Wei et al., 2021)	69.4±0.35 / 80.1±0.41	52.4±0.32 / 70.3±0.28	42.9±1.45 / 67.4±1.07
FixMatch + SAW	**78.7**±0.77 / **84.2**±0.36	**64.3**±1.96 / **76.4**±0.88	**57.5**±2.83 / **70.5**±1.50

Lai, Z., C. Wang, S.-C. Cheung, and C.-N. Chuah. 2022. SaR: Self-adaptive Refinement on Pseudo Labels for Multiclass-Imbalanced Semi-supervised Learning. In Workshop on Learning with Limited Labelled Data for Image and Video Understanding @ CVPR

Lai, Z., C. Wang, L.C. Oliveira, B. Dugger, S.-C. Cheung, and C.-N. Chuah. 2022. Smoothed Adaptive Weighting for Imbalanced Semi-Supervised Learning: Improve Reliability Against Unknown Distribution Data. In ICML 2022

Results

SSL: Using only 0.1% regions as label set

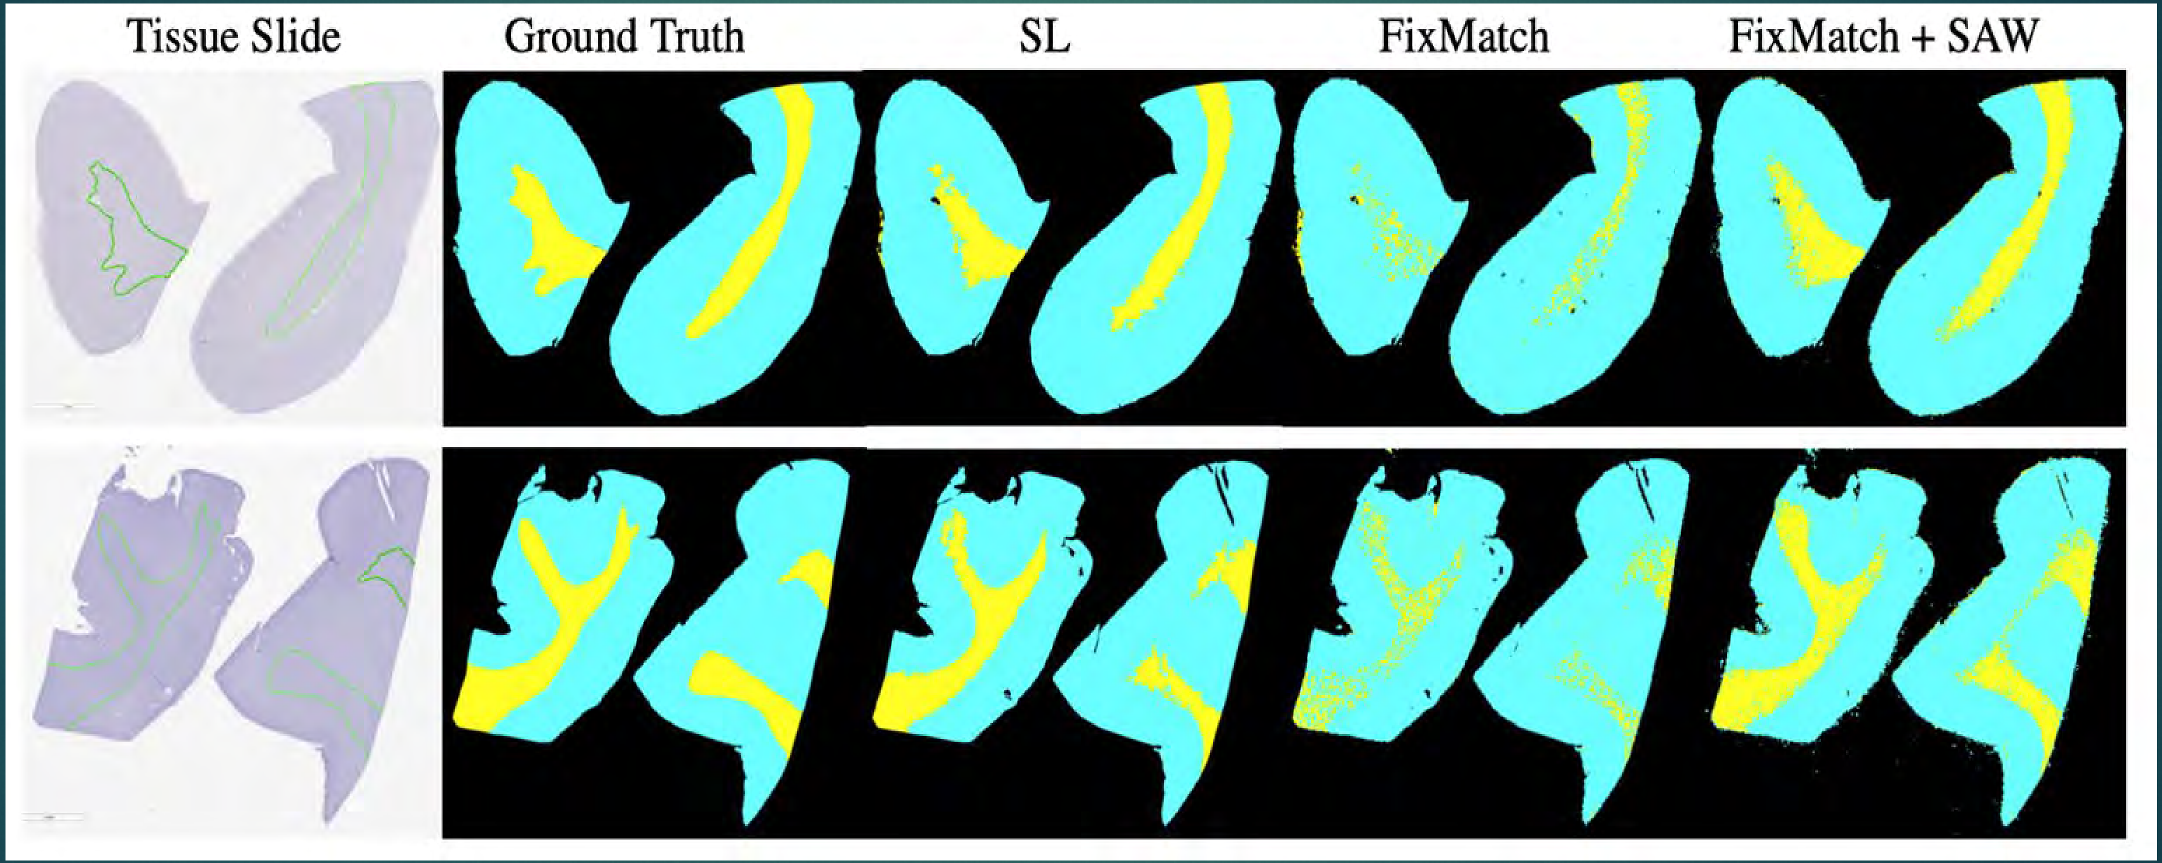

Active Learning (AL)

- Identify a small number of highly informative unlabeled data for expert labeling

- Active Learning
  1. Identify the most informative samples S_t with respect to classifier C_{t-1}
  2. Send to expert labeler to get labels L_t
  3. Improve $C_{t-1} \rightarrow C_t$ with (S_t, L_t)
  4. $t + 1 \rightarrow t$ and repeat

- Measure informativeness via an Acquisition Function
  - Uncertainty of the current classifier
  - Diversity of samples

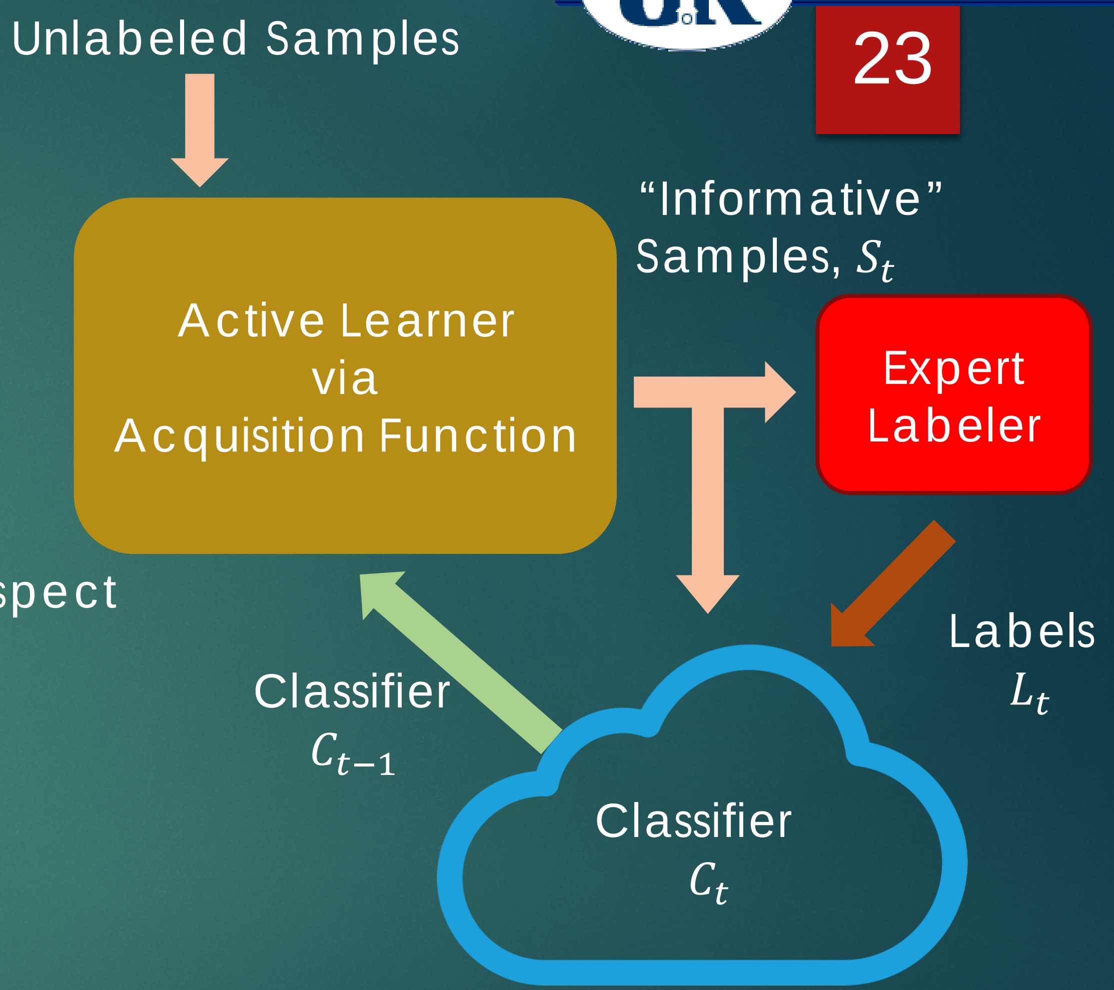

Combining AL and SL

- Active Learning
  - Cold-start problem: limited starting set may result in high-biased selection
  - High computational complexity

- Semi-supervised learning
  - Relieve the cold-start problem in AL by minimizing confirmation bias
  - Reduce AL complexity by using pseudo labeling to identify uncertain samples

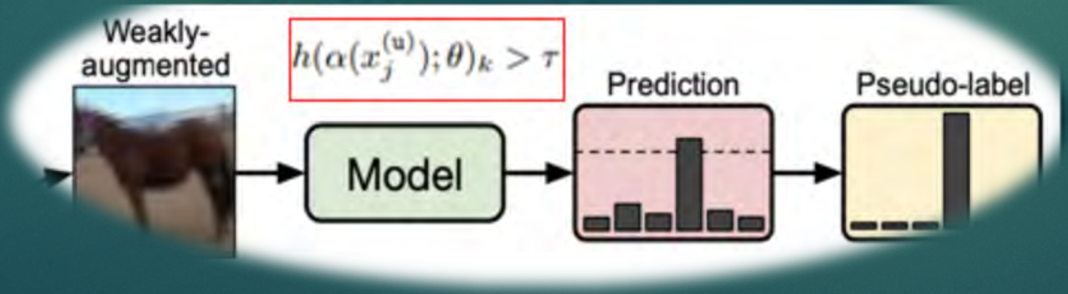

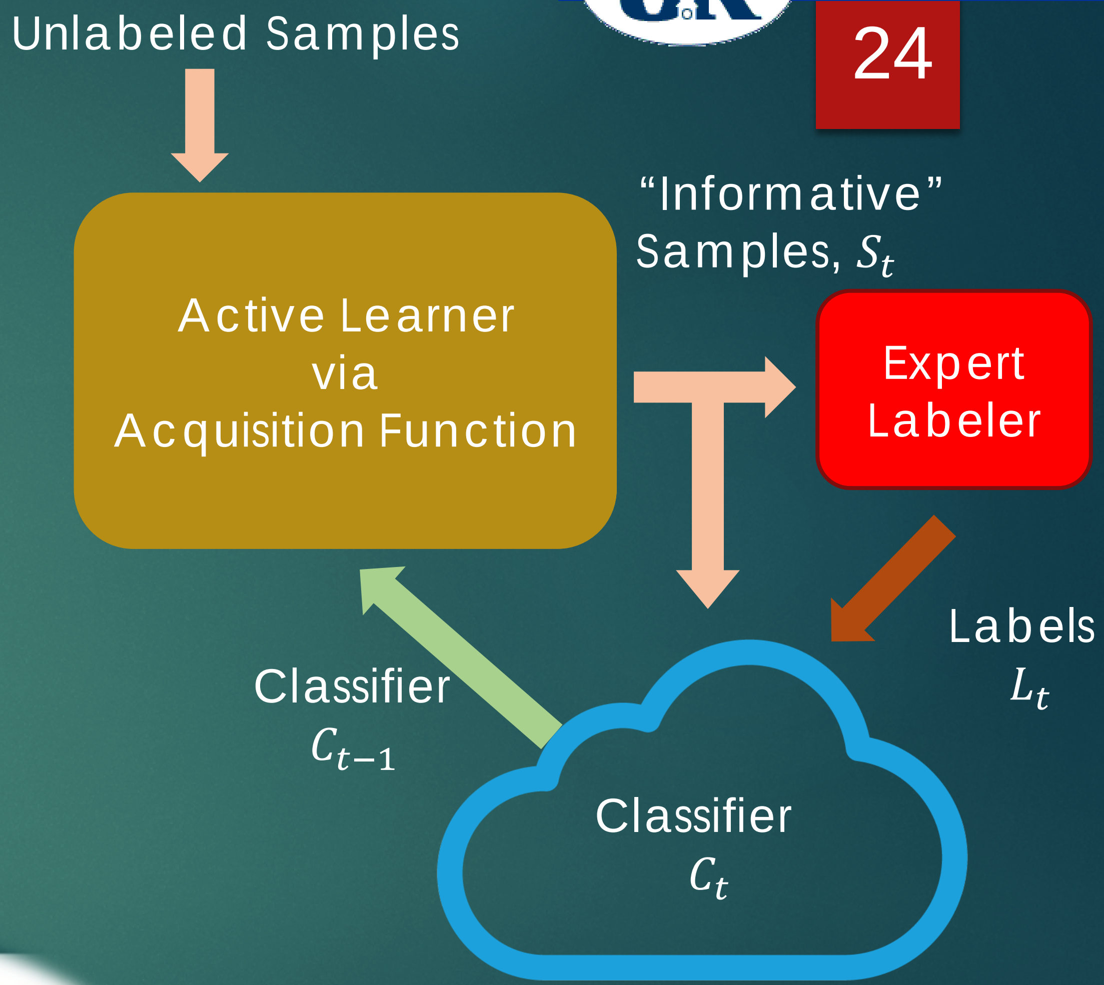

Combining AL and SL for WSI

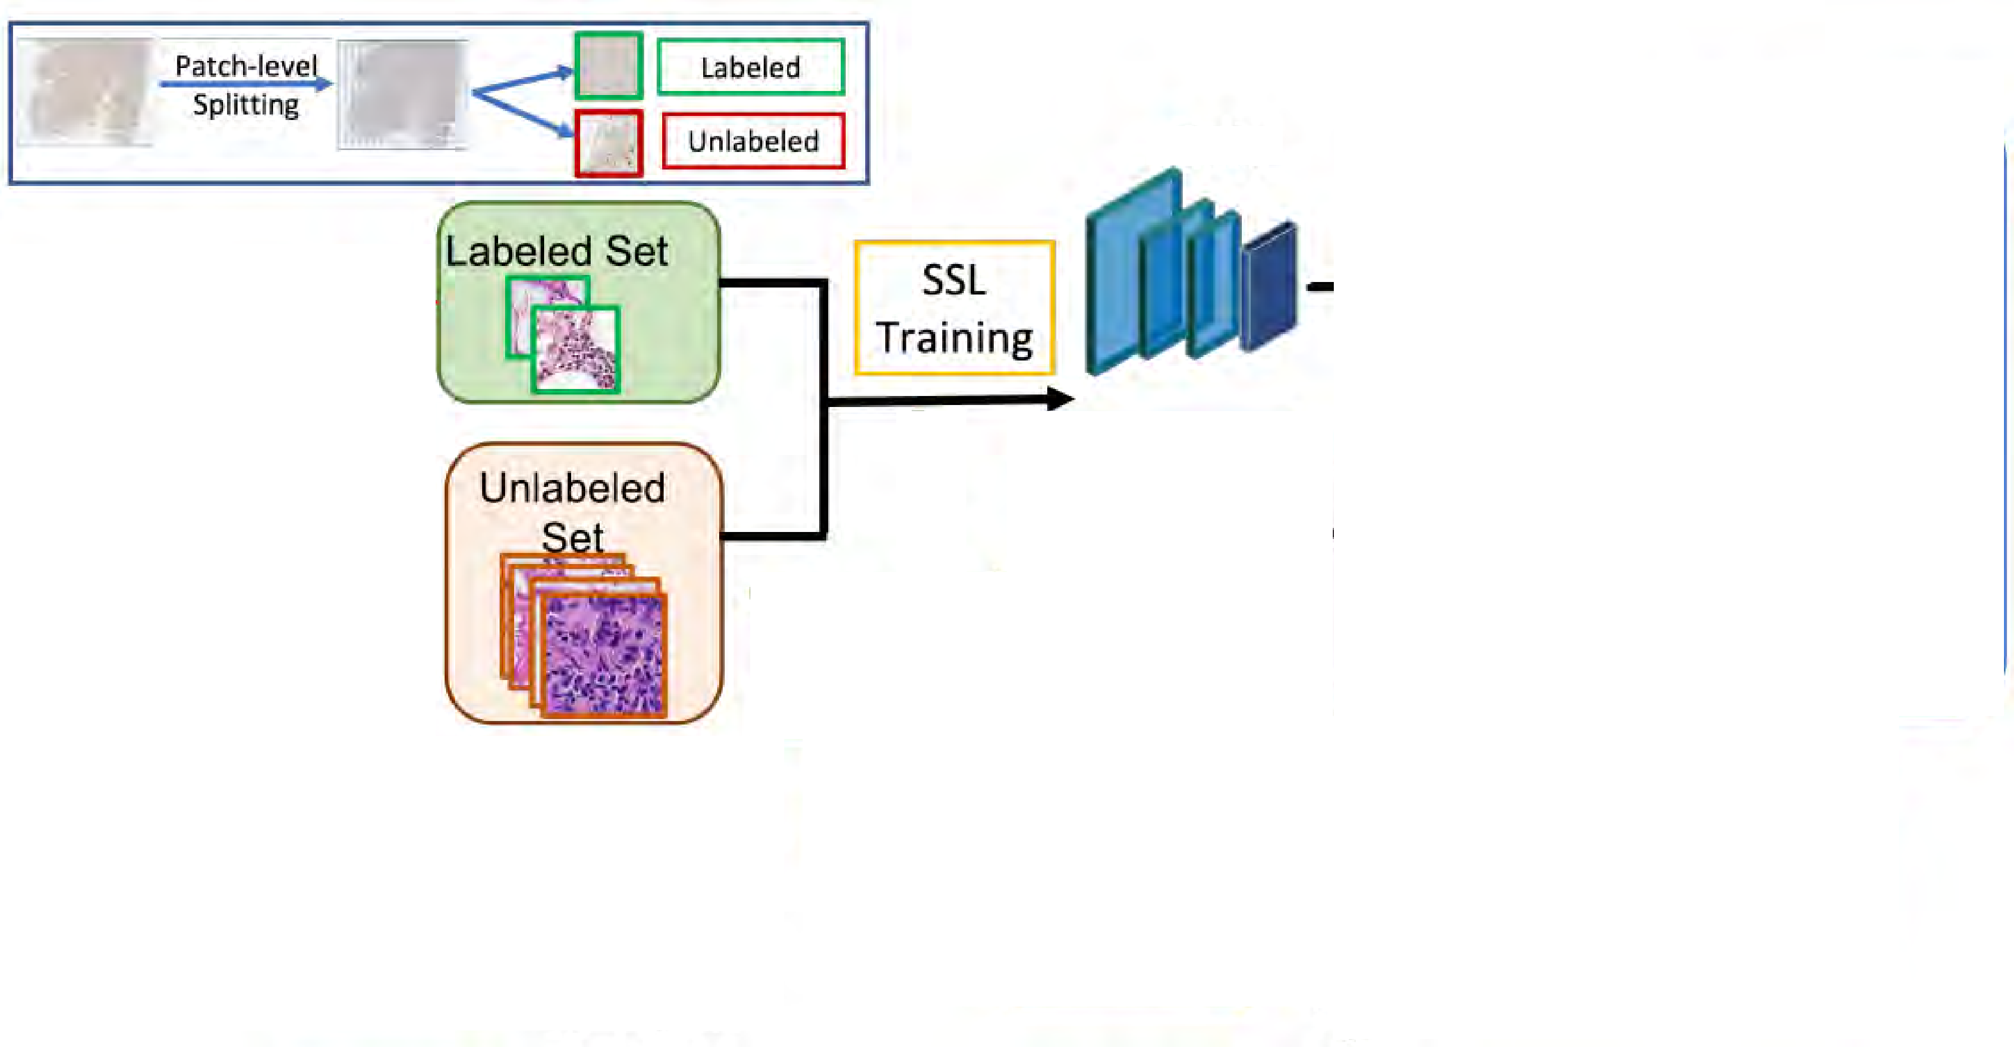

Results

PIXEL-WISE IOU SCORES FOR AD, NAD, AND OVERALL TEST SET

Method	FCN [6]		U-Net [7]		FixMatch [15]	ICCVW [29]	SemA-Path
Labeled data	**2 WSIs**	**All WSIs**	**2 WSIs**	**All WSIs**	**0.1%**	**0.1%**	**0.1%**
AD Back	61.04 ± 5.44	81.13 ± 9.17	59.74 ± 13.9	**96.80** ± 1.48	93.15 ± 2.41	95.01 ± 1.17	95.09 ± 1.21
AD GM	46.98 ± 2.78	76.07 ± 8.91	37.16 ± 9.93	**89.58** ± 5.12	78.57 ± 3.87	88.80 ± 3.92	88.91 ± 4.05
AD WM	27.75 ± 5.50	62.23 ± 14.0	7.57 ± 6.02	**82.53** ± 7.70	56.66 ± 16.4	81.83 ± 5.53	81.95 ± 4.58
NAD Back	66.66 ± 5.17	88.42 ± 1.55	78.46 ± 18.5	**97.36** ± 3.15	97.07 ± 0.31	97.26 ± 0.52	97.33 ± 0.78
NAD GM	50.15 ± 0.49	79.37 ± 2.95	59.59 ± 13.6	**94.42** ± 3.30	83.97 ± 7.76	93.47 ± 1.60	93.59 ± 1.55
NAD WM	19.72 ± 13.6	49.89 ± 12.8	3.02 ± 3.09	**81.25** ± 9.53	22.72 ± 19.0	75.85 ± 11.4	77.95 ± 10.9
Background	63.29 ± 5.81	84.05 ± 9.17	68.28 ± 17.2	**97.02** ± 2.15	94.72 ± 2.71	95.91 ± 1.48	95.99 ± 1.33
GM	48.25 ± 2.66	77.39 ± 7.06	46.13 ± 15.8	**91.52** ± 4.94	80.73 ± 6.01	90.67 ± 3.90	90.78 ± 3.34
WM	24.54 ± 9.80	57.29 ± 14.3	5.75 ± 5.37	**82.02** ± 7.98	43.08 ± 24.0	79.44 ± 8.34	80.35 ± 8.67
Mean	45.36 ± 3.26	72.91 ± 7.56	40.05 ± 10.2	**90.19** ± 3.84	72.84 ± 7.18	88.67 ± 3.12	89.04 ± 2.99

The results are from the hold-out test set. AD refers to Alzheimer's disease cases while NAD refers to Non-Alzheimer's disease cases. 2 WSIs refers to 2 WSIs are labeled, equivalent to 10% regions of all WSIs; all WSIs refers to all WSIs are labeled. 0.1% refers to 0.1% regions of all WSIs are labeled, which can be tiled into 600 patches; so as 0.07% which can be tiled into 400 patches.

Lai, Zhengfeng, J. Chauhan, D. Chen, B. N. Dugger, S.-C. Cheung & C.-N. Chuah. 2024. Semi-Path: An Interactive Semi-Supervised Learning Framework for Gigapixel Pathology Image Analysis. Smart Health 32 (100474): 100474.

Results

Both SSL results use FixMatch as the backbone and use 0.1% labeled area of 20 WSIs in the training set. SemA-Path uses 3 AL cycles to get to 0.1%.

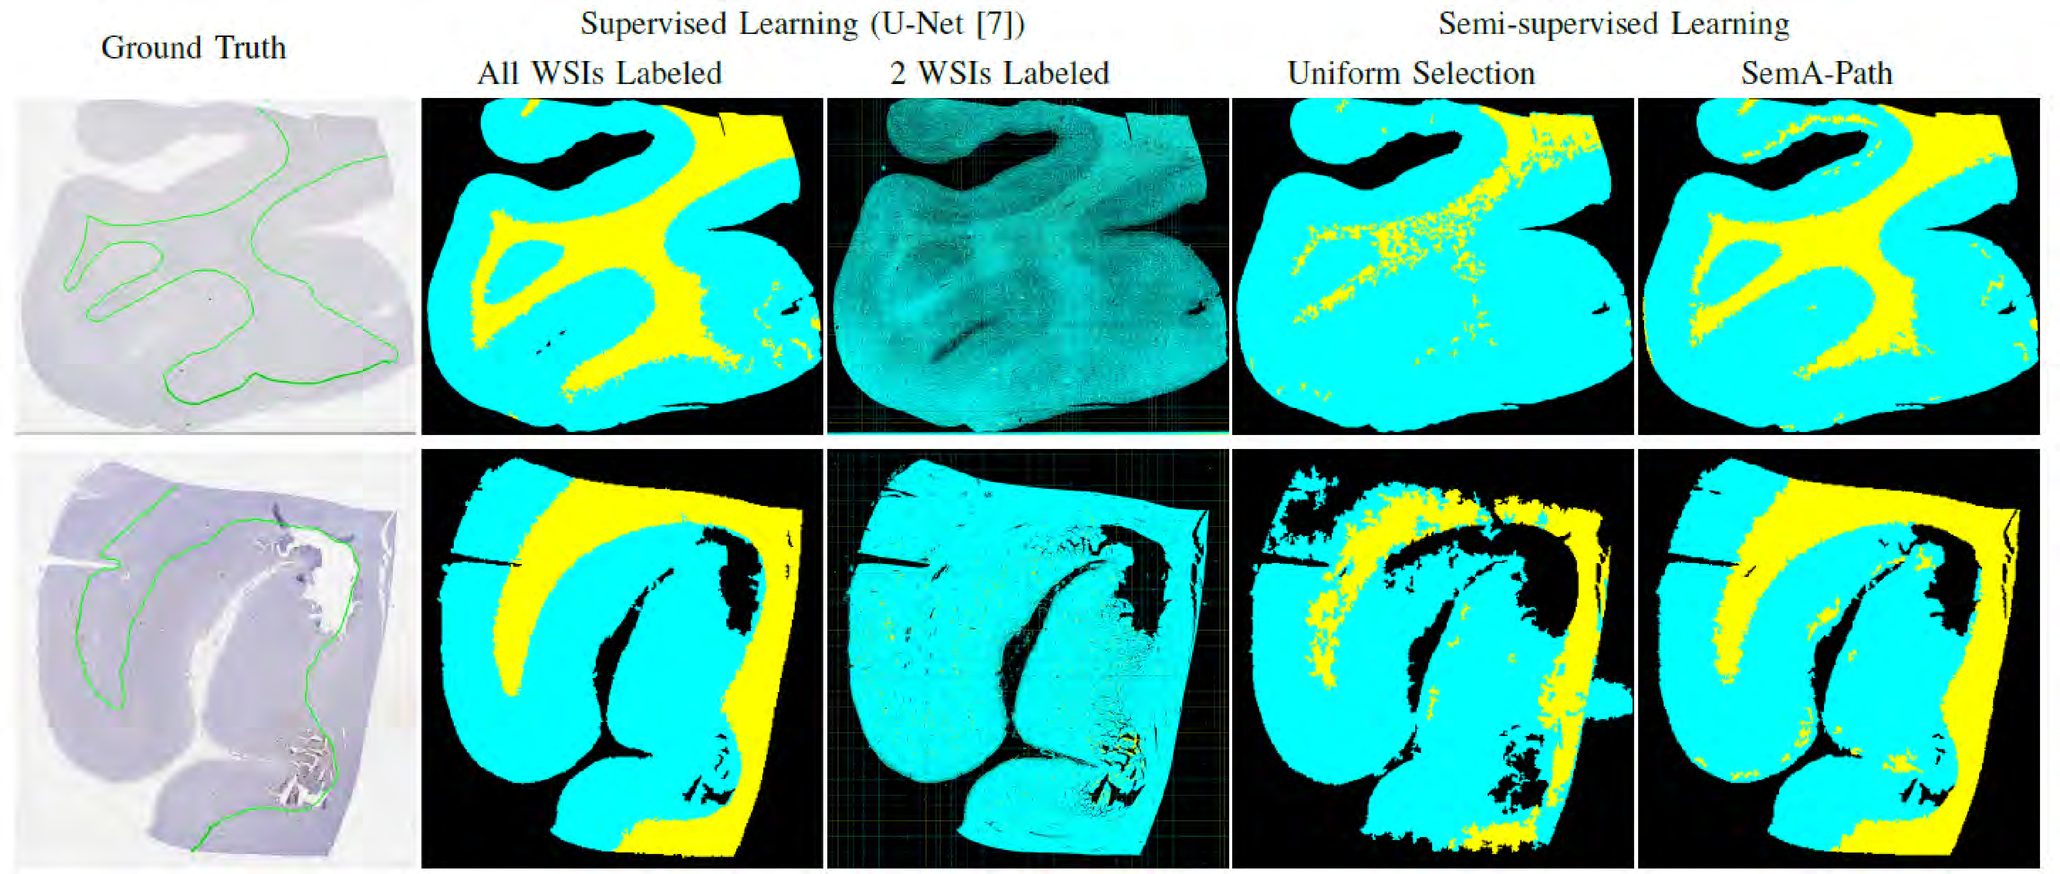

Lai, Z., C. Wang, L. C. Oliveira, B. N. Dugger, S. C. Cheung, and C. N. Chuah. 2022. SemA-Path: Semi-supervised Active Learning with Inner-Outer Selection for Pathology Image Classification and Segmentation. Submitted to IEEE Trans. On Medical Imaging.

Data Privacy in Distributed Learning

UNIVERSITY OF KENTUCKY

Privacy Challenge in AI

- Raw data contain sensitive information (e.g. health, intelligence, financial, etc.) and cannot leave the premise
- Local site users may not trust cloud (same cloud may also serve competitors)
- Traditional end-to-end encryption only protects storage and transfer, not calculations
- One of the top problems in AI system challenges

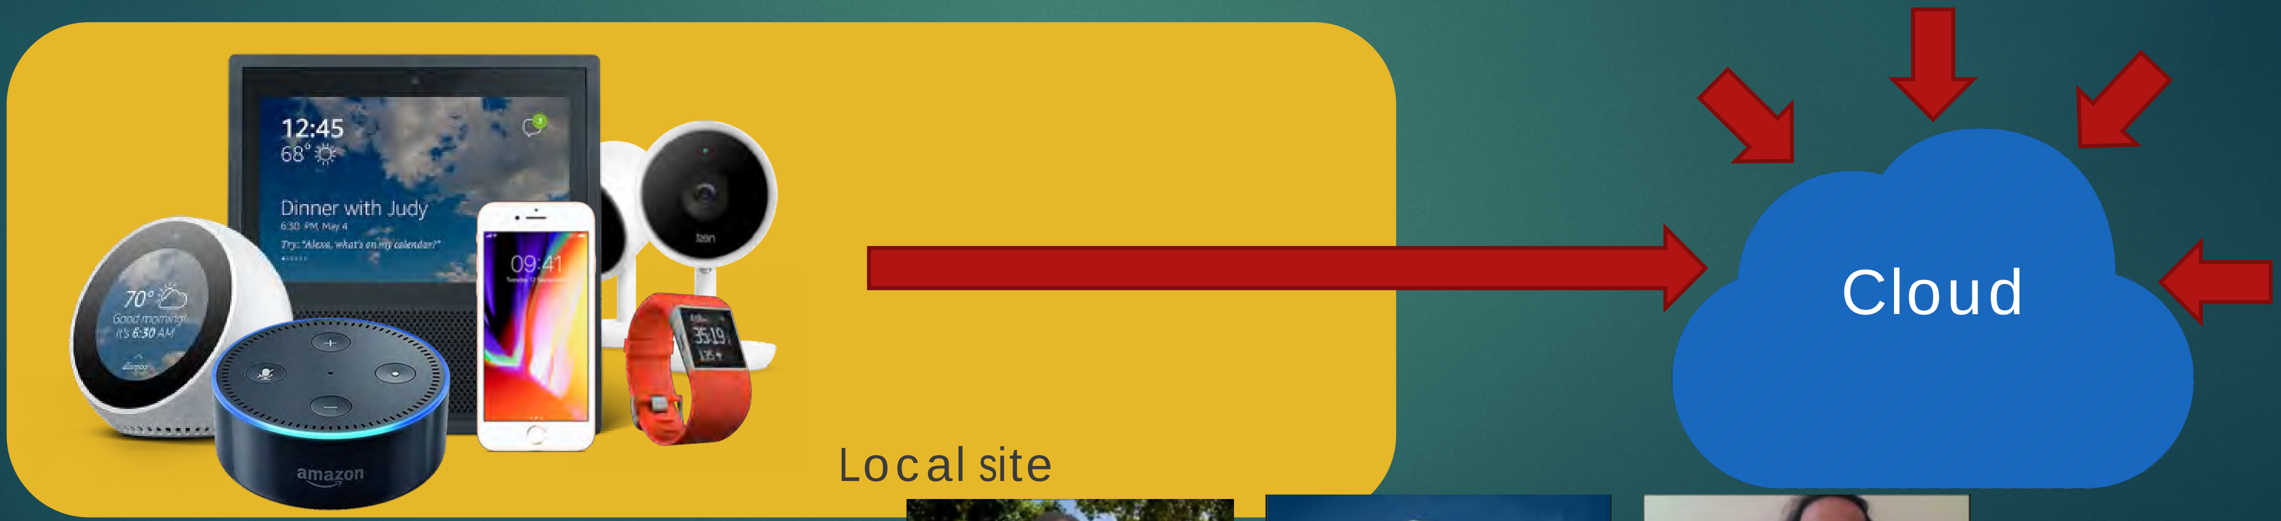

PPML Approaches

- Redaction
- Federated Learning
- Encrypted-domain processing
- Differential Privacy
- Synthetic Data

DP Synthetic Models

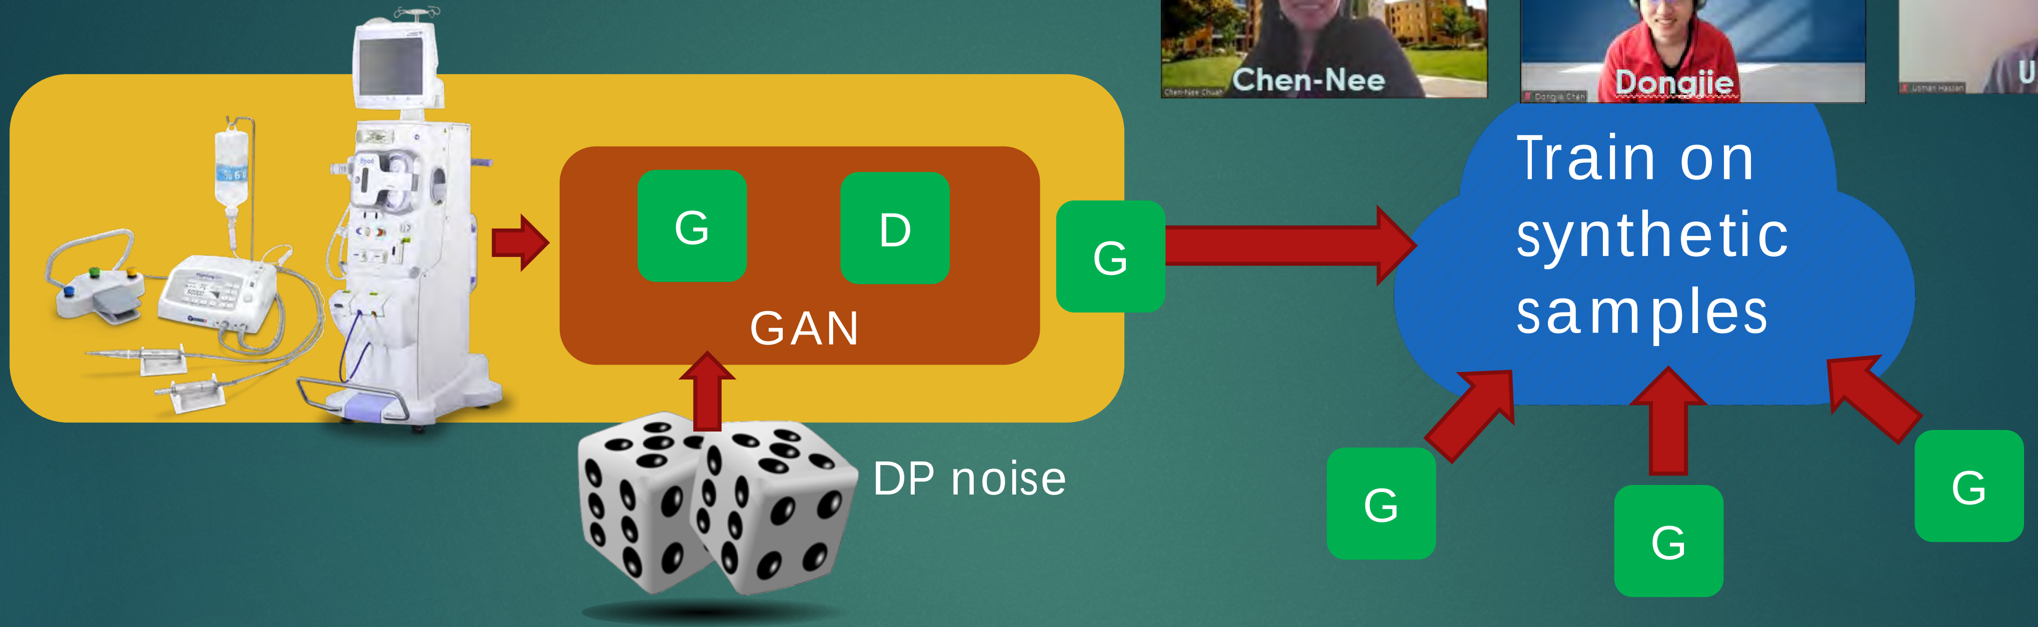

- Use GAN trained on sensitive data to generate synthetic surrogate
- Use **Differential Privacy** in the training of GAN to protect private data
- Main advantages over other PPMLs:
  - No changes on any downstream ML tasks
  - Support human-in-the-loop operations such as active learning
- Potential limitations : poorer quality than real images

Differential Privacy

- Perturb output to make it “almost indistinguishable” when run with or without any sample (neighboring)

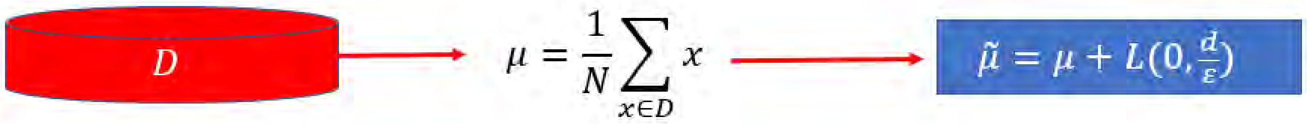

Laplacian/Gaussian Mechanism

Differential Privacy

- Perturb output to make it "almost indistinguishable" when run with or without any sample (neighboring)
- Privacy budget ϵ : smaller means more privacy but poorer quality
- Definition: For any neighboring datasets D_1 and D_2, we have

$$P[A(D_1) \in S] \leq e^{\varepsilon} \cdot P[A(D_2) \in S] \text{ for all } S \text{ in Range(A)}$$

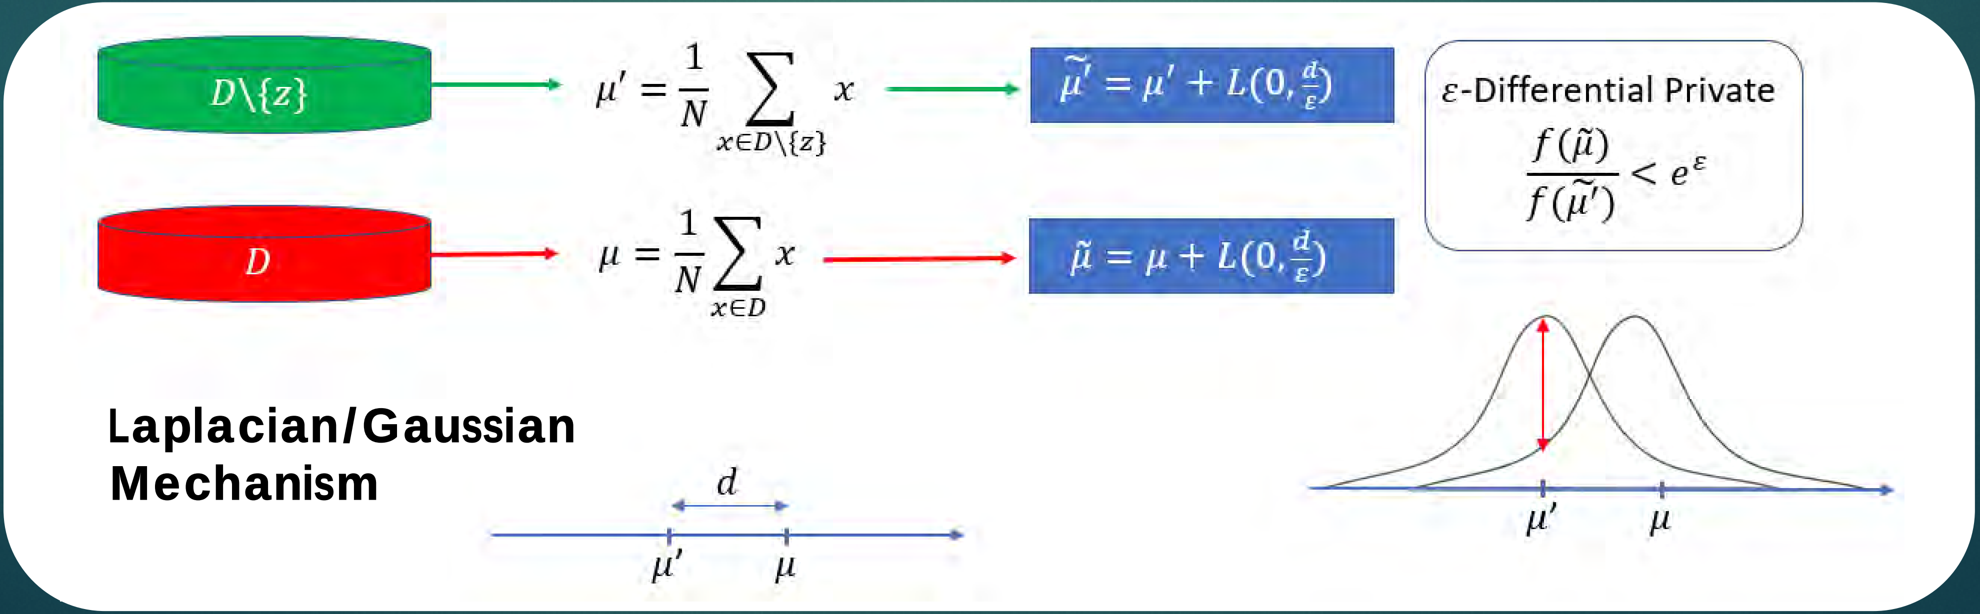

Adding DP to deep learning

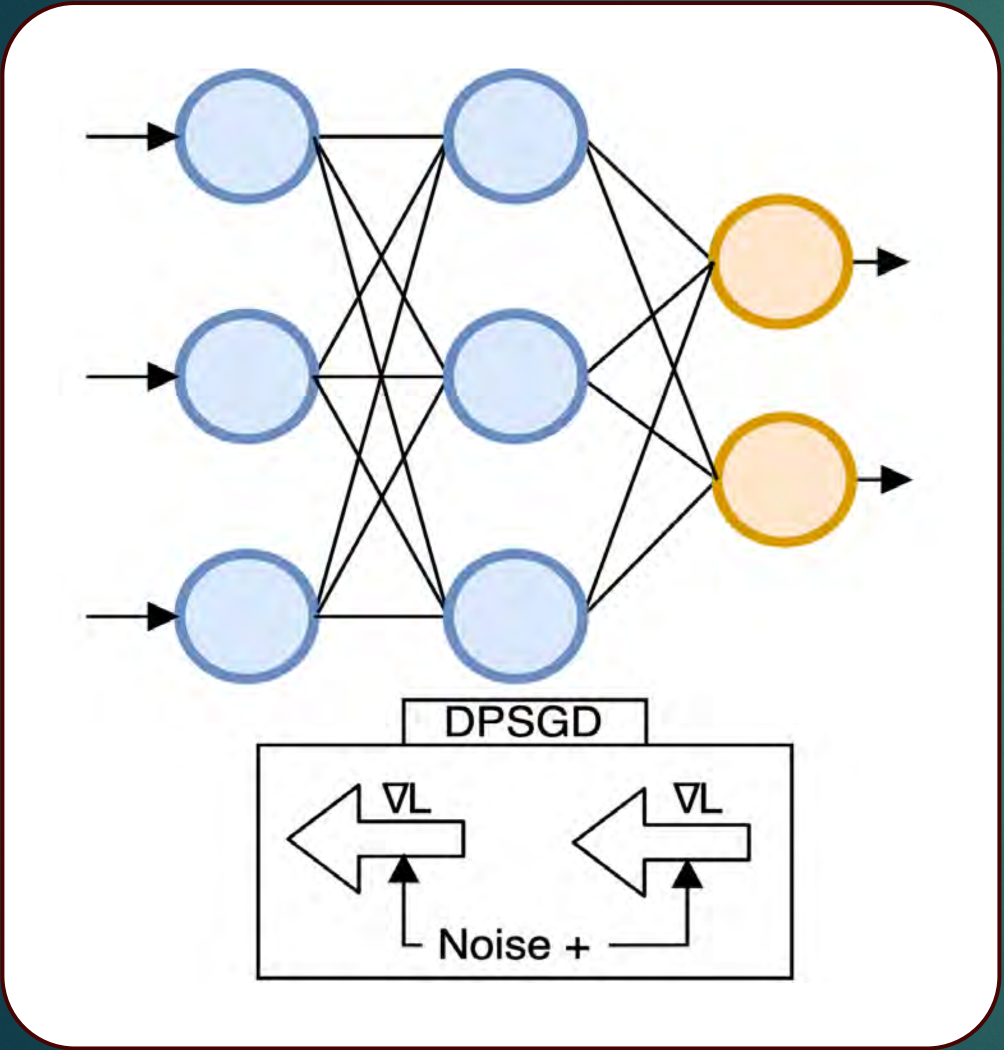

Problem with DPSGD + GAN : non-convergence or converge to a noisy equilibrium.

Adding DP to deep learning

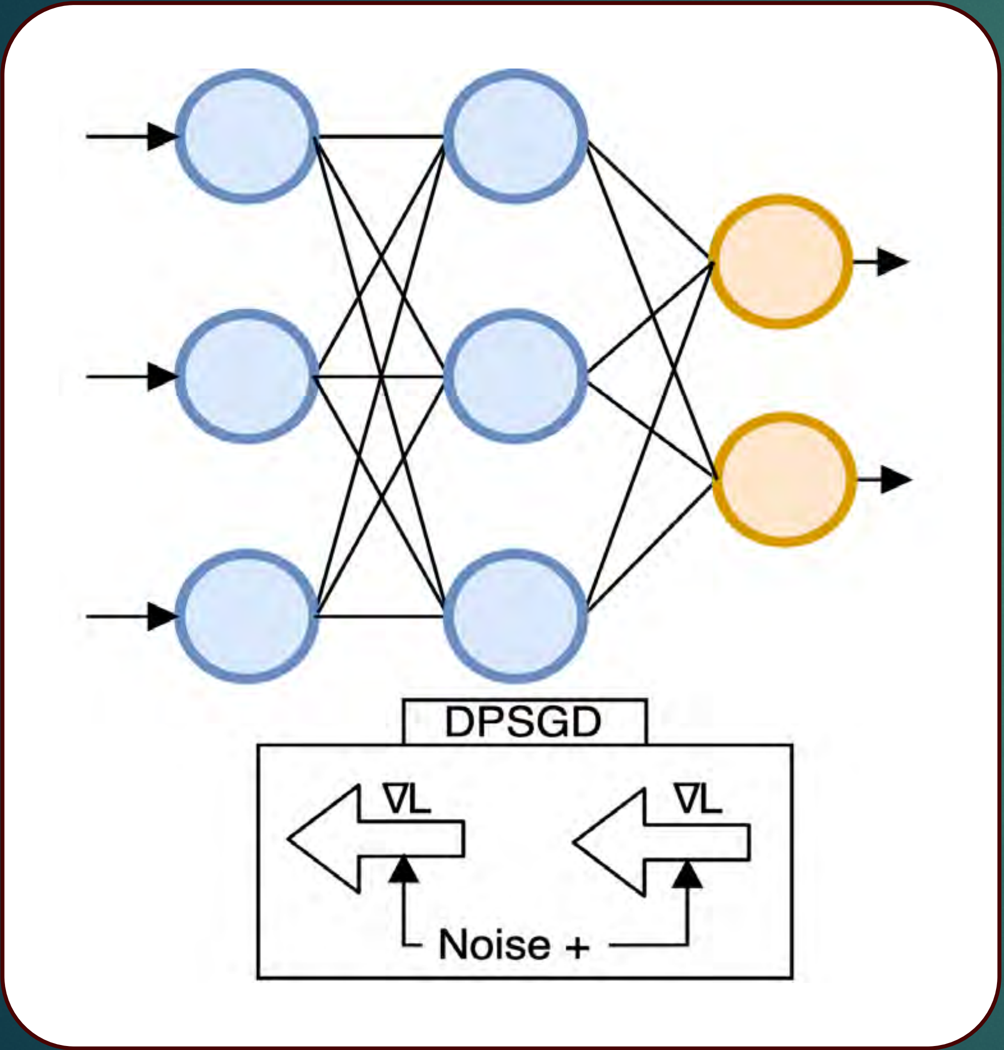

Problem with DPSGD + GAN : non-convergence or converge to a noisy equilibrium.

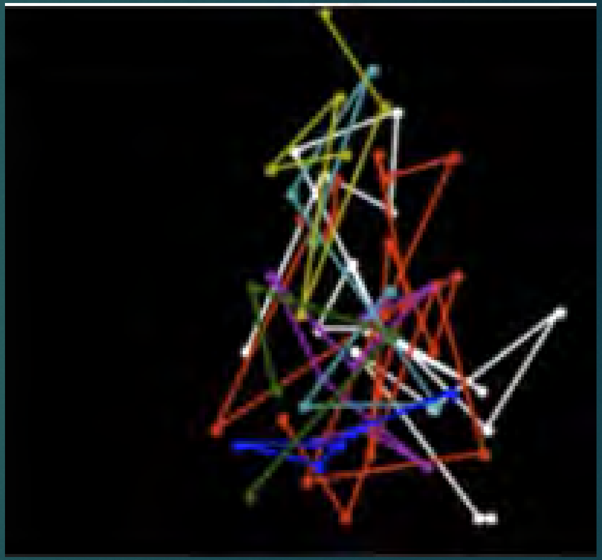

DP Latent-GANs

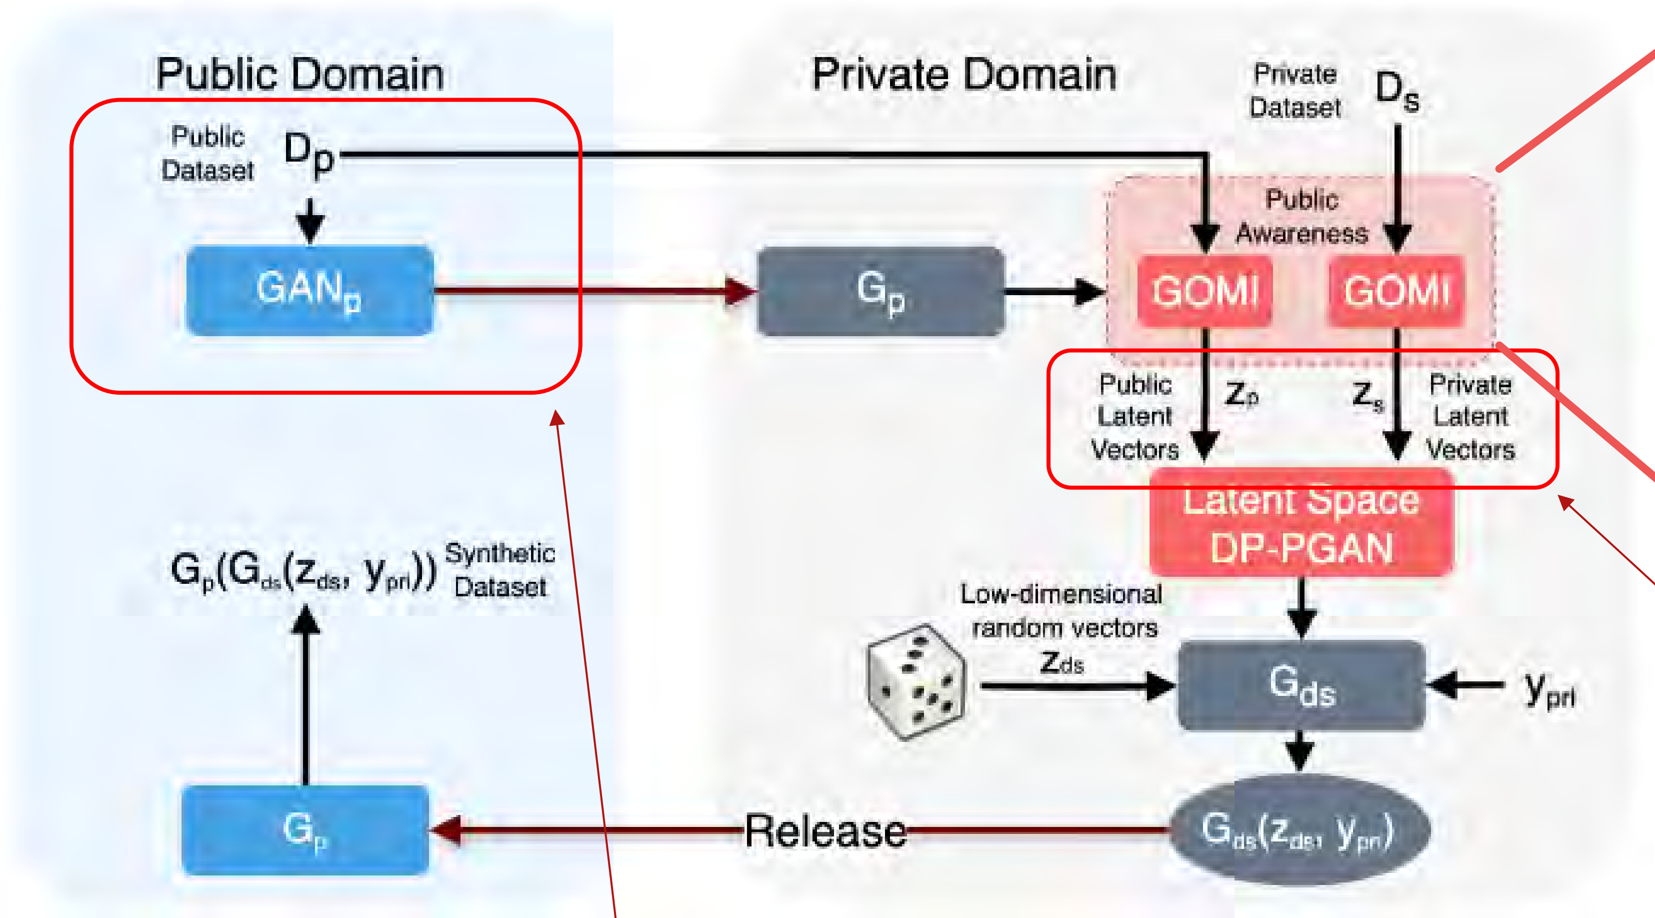

Differentially Private Publicly-trained Adversarial Model Inversion (DP-PAMI)

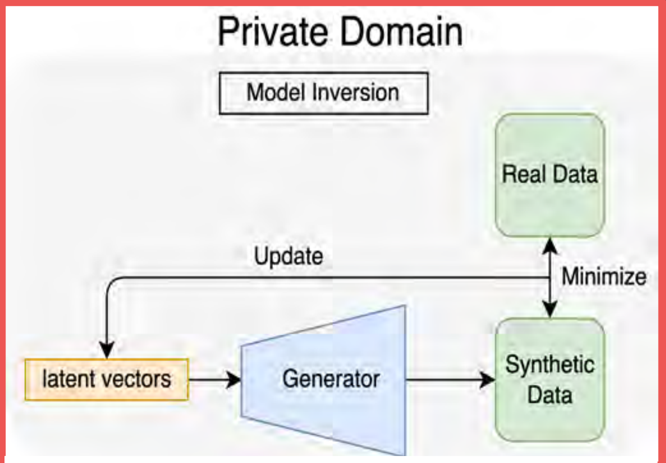

$$z_s = \arg\min_z \quad ||G_p(z) - x_s||^2 / P(z)$$

Lower dimension and Gaussian regularization make Latent-GAN easier to train with DP Noise

Use of publicly trained GAN to build basic image generation

Results

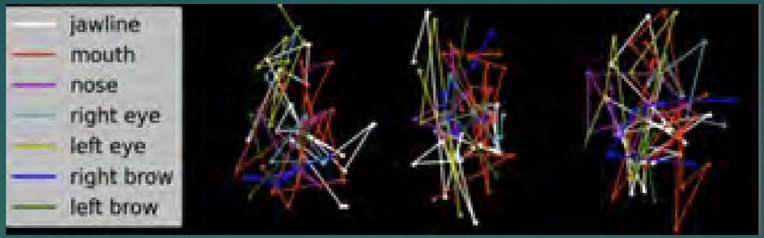

DP-GAN (ε=10)

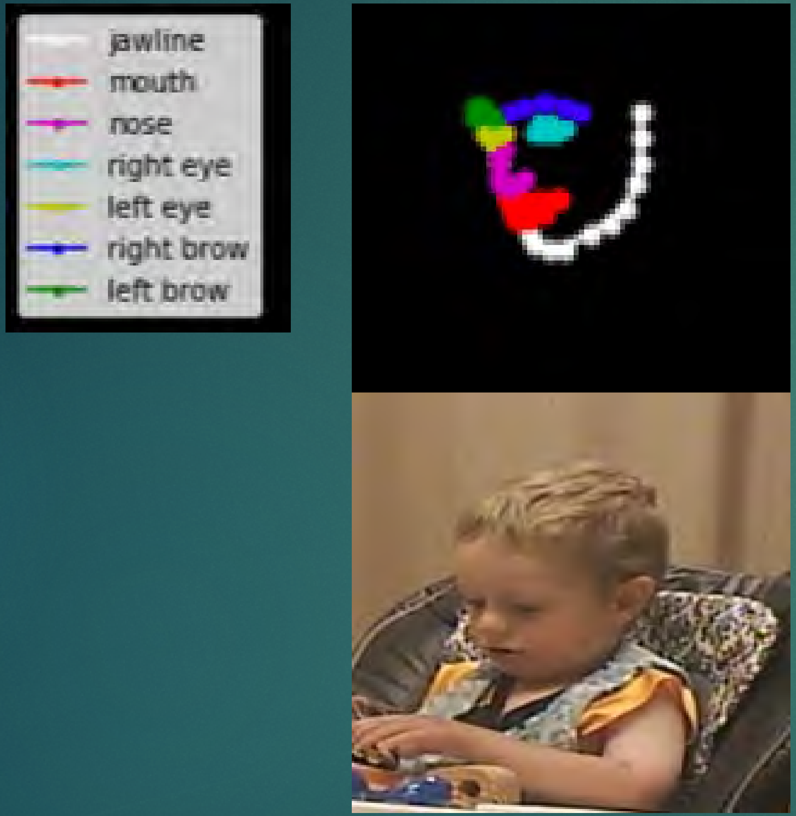

Facial Landmarks from human subjects

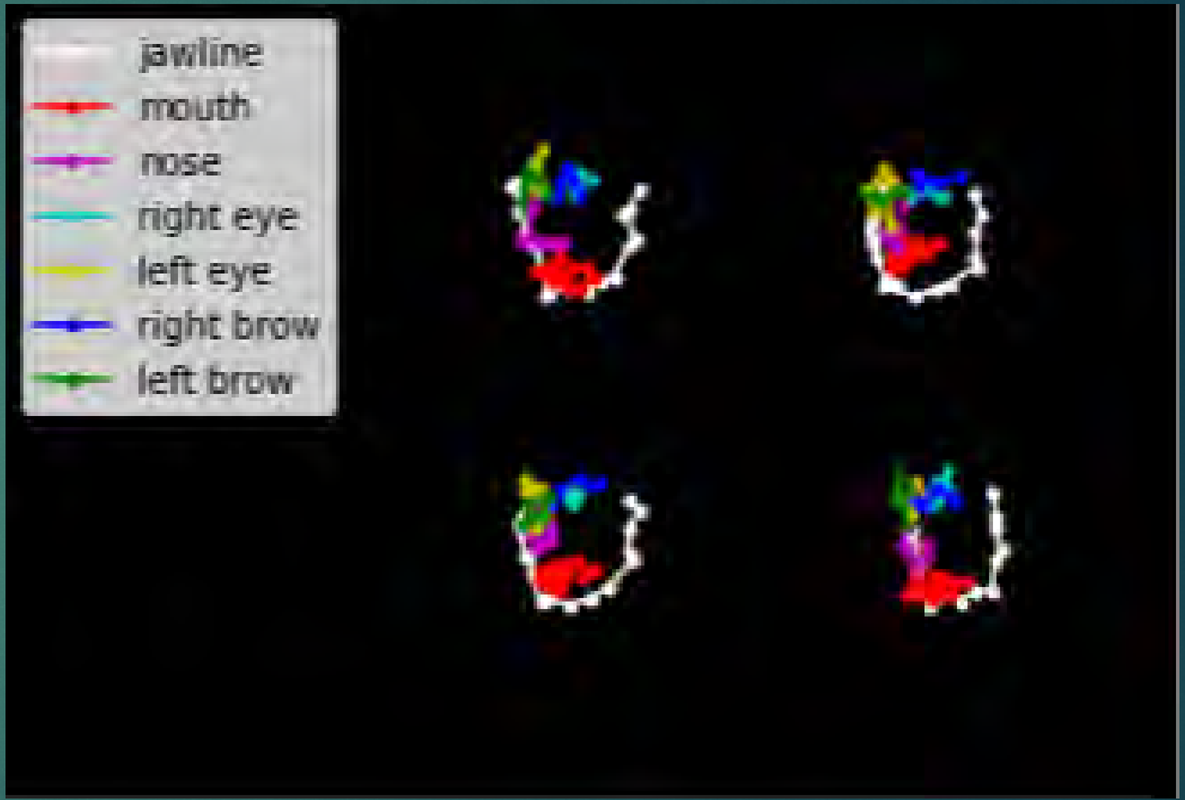

Synthetic Facial Landmarks from DPMI-GAN (ε=10)

Dongjie Chen, et al. 2024. DP-PAMI: A Latent Space Solution for Differentially Private Synthetic Data Generation. Submitted to IEEE Transactions on Information Forensics and Security.

Results

TABLE I: FID values for different methods and ϵ values for CIFAR10 and SVHN.

Method	CIFAR10 $\epsilon=1$	CIFAR10 $\epsilon=5$	CIFAR10 $\epsilon=10$	CIFAR10 $\epsilon=20$	CIFAR10 $\epsilon=50$	SVHN $\epsilon=1$	SVHN $\epsilon=5$	SVHN $\epsilon=10$	SVHN $\epsilon=20$	SVHN $\epsilon=50$
DP-GAN [31]	323.27	329.80	336.21	255.29	247.40	306.54	295.11	297.73	290.70	253.32
DP-MERF [38]	331.28	325.04	324.78	312.54	307.72	344.58	338.22	327.84	320.06	310.43
G-PATE [37] (cond.)	444.56	439.19	347.86	309.03	309.03	461.00	416.02	461.08	402.51	400.57
GS-WGAN [35] (cond.)	354.46	275.70	233.30	223.62	223.62	302.13	158.38	162.19	161.48	119.4
DP-MEPF [24] (ϕ_1, ϕ_2)	175.50	166.88	151.48	152.24	152.91	113.54	95.91	120.22	115.90	87.15
DP-MEPF [24] (ϕ_1)	132.57	128.92	124.01	111.99	104.98	101.05	93.16	82.60	81.76	78.69
DPMI [23]	130.61	121.67	108.06	104.47	97.68	72.27	83.96	72.67	67.91	63.62
DPGOMI [25]	127.67	95.54	94.45	93.67	93.14	70.13	67.47	65.47	55.64	53.88
DP-PAMI	**108.0**	**92.88**	**86.90**	**86.79**	**86.32**	**66.87**	**63.19**	**56.80**	**55.11**	**48.22**
DP-PAMI (ϵ=∞)	79.46	79.46	79.46	79.46	79.46	40.45	40.45	40.45	40.45	40.45

TABLE II: Inception Score and Downstream Classification Precision comparison on $\epsilon = 10$

	Inception Score CIFAR10	Inception Score SVHN	Classification CIFAR10	Classification SVHN
DP-GAN [31]	1.67	1.73	0.28	0.32
G-PATE [37] (cond.)	1.29	1.46	0.32	0.45
GS-WGAN [35] (cond.)	1.88	1.63	0.31	0.35
DP-MERP [38]	2.95	2.39	0.35	0.53
DP-MEPF [24] (ϕ_1, ϕ_2)	3.05	2.44	0.67	0.67
DP-MEPF [24] (ϕ_1)	2.97	2.61	0.71	0.77
DPMI [23]	4.46	2.07	0.67	0.69
DPGOMI [25]	4.74	2.59	0.73	0.79
DP-PAMI	**4.81**	**2.67**	**0.80**	**0.81**
DP-PAMI (ϵ=∞)	5.02	2.76	0.87	0.92

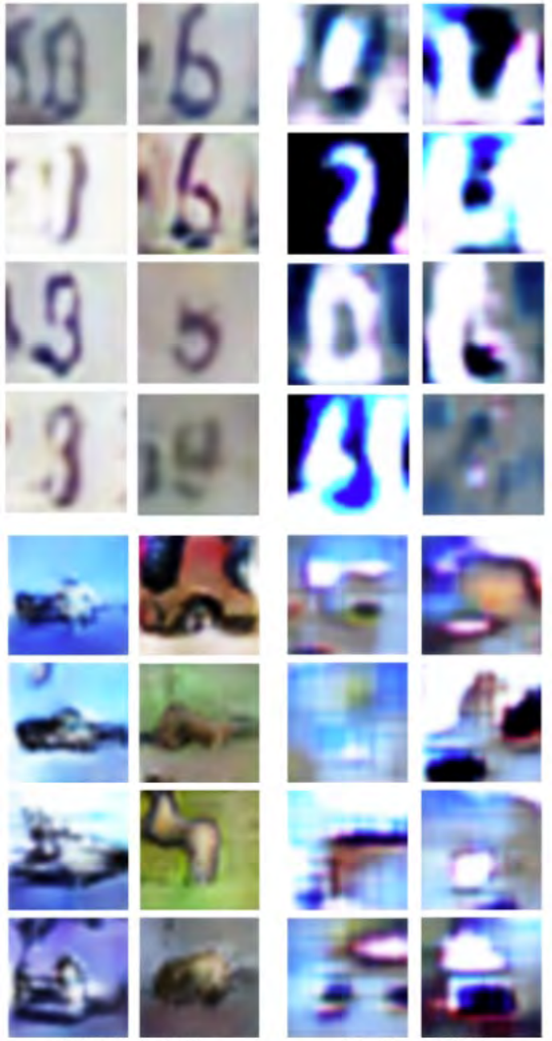

(a) DP-PAMI (b) DP-MEPF

Visual comparison at ϵ = 50

Exponential Mechanism

$D\backslash\{z\}$ → $\mu' = \frac{1}{N}\sum_{x\in D\backslash\{z\}} x$ → $\widetilde{\mu'} = \mu' + L(0, \frac{d}{\varepsilon})$

D → $\mu = \frac{1}{N}\sum_{x\in D} x$ → $\tilde{\mu} = \mu + L(0, \frac{d}{\varepsilon})$

ε-Differential Private

$$\frac{f(\tilde{\mu})}{f(\widetilde{\mu'})} < e^{\varepsilon}$$

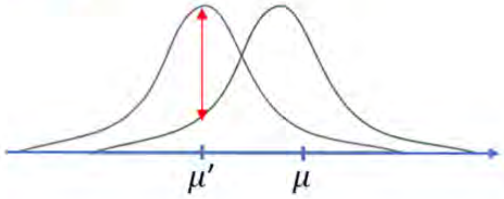

Laplacian/Gaussian Mechanism

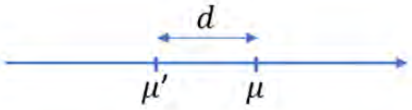

Use EM to obfuscate the distribution of the private latent vectors

$D\backslash\{z\}$ → $\mu' = \frac{1}{N}\sum_{x\in D\backslash\{z\}} x$ → $f(z) \sim \exp(-\epsilon\, U(z, \mu'))$

D → $\mu = \frac{1}{N}\sum_{x\in D} x$ → $f(z) \sim \exp(-\epsilon\, U(z, \mu))$

ε-Differential Private

$$\frac{f(\tilde{\mu})}{f(\widetilde{\mu'})} < e^{\varepsilon}$$

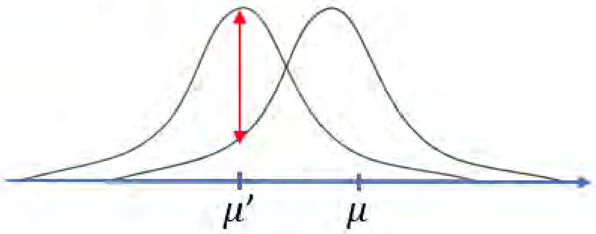

Exponential Mechanism

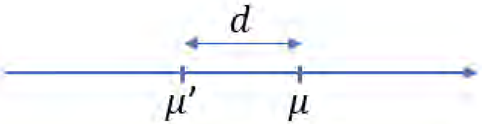

DPGEM

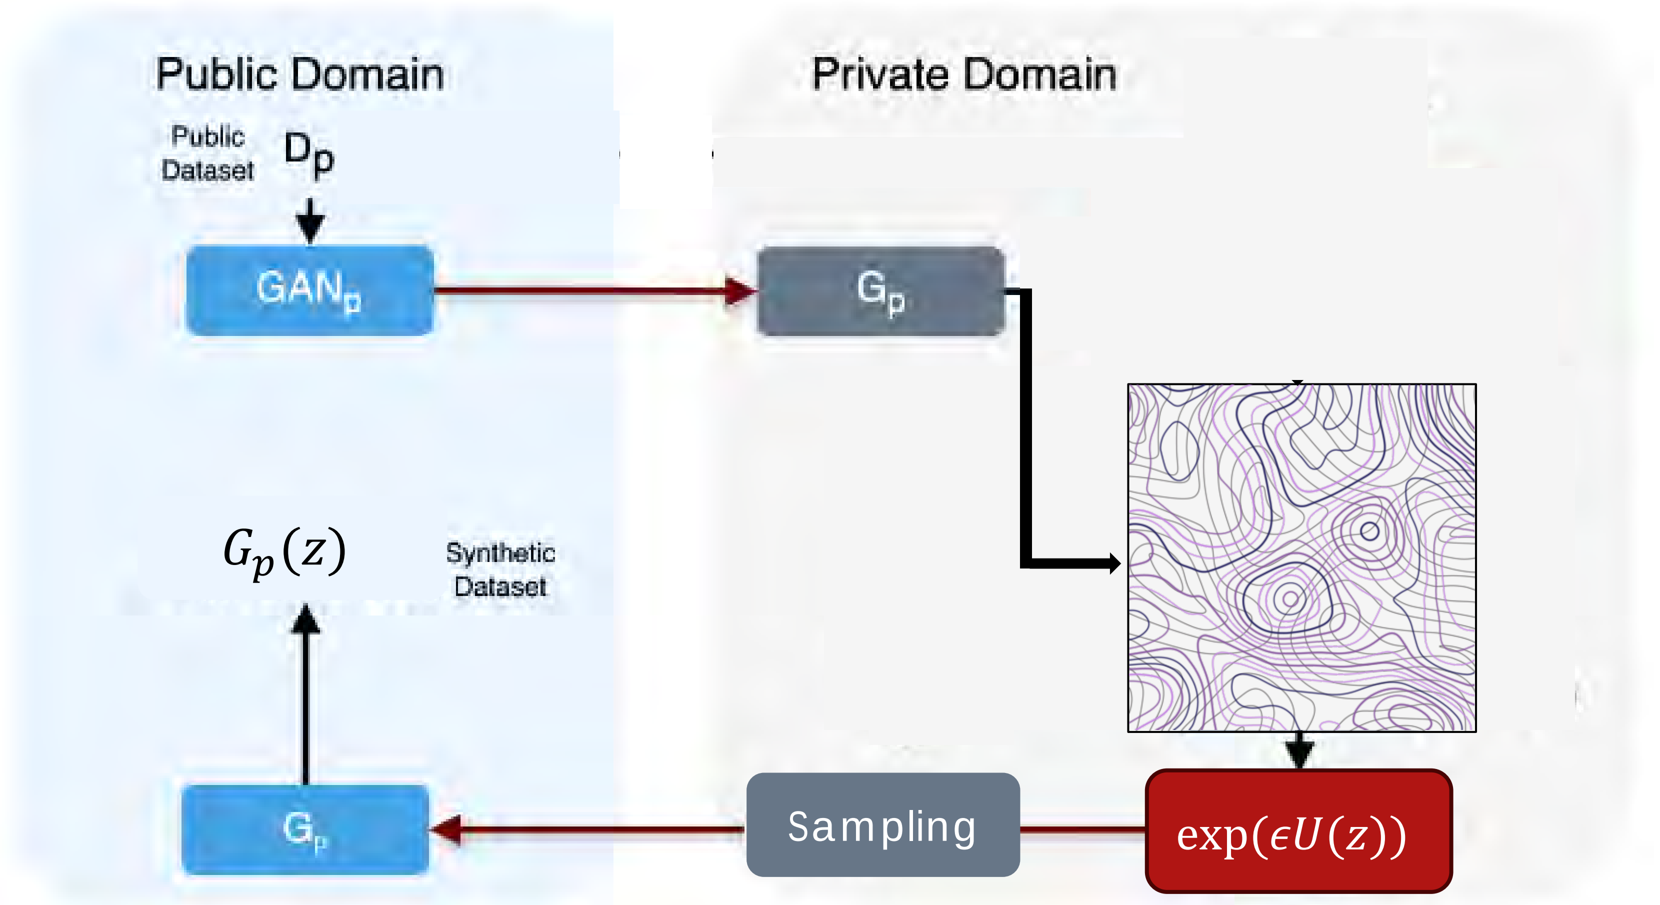

Differentially Private Generative Model with Exponential Mechanism (DPGEM)

Sampling in latent space is still changing

- Gradient information is readily available
- Gradient-based sampling method: Hamiltonian Monte Carlo

DPGEM

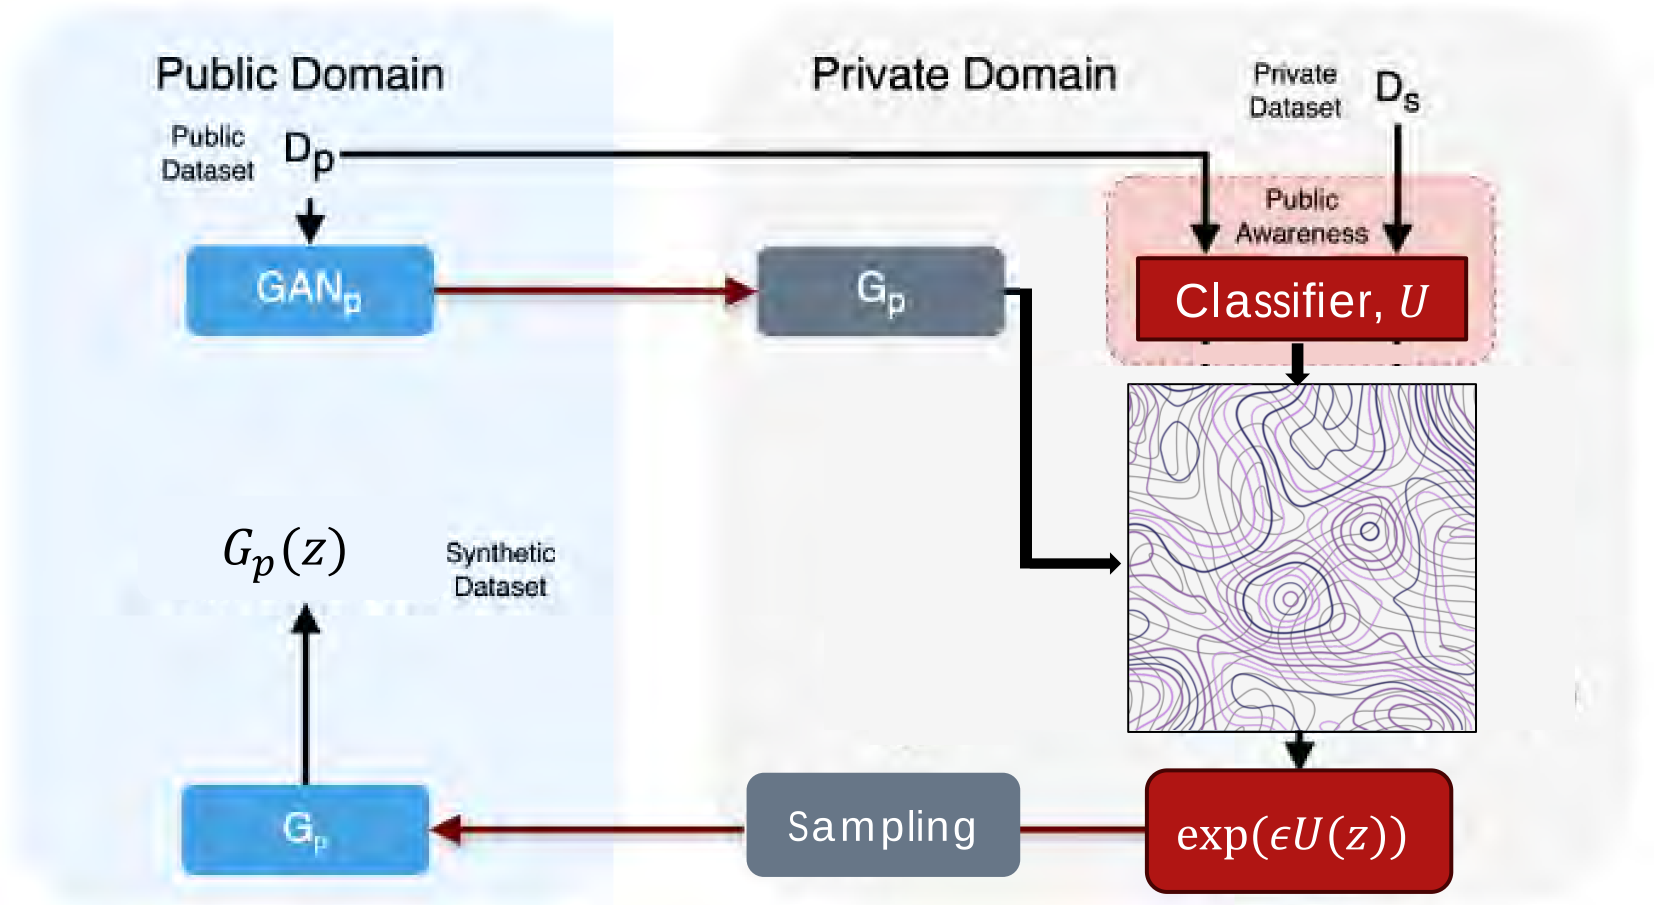

Differentially Private Generative Model with Exponential Mechanism (DPGEM)

Limitation

- Every synthetic sample reveals sensitive private information
- Unlike DP-GAN which poses no limit on #'s of synthetic samples
- Need clever Privacy Accounting method

DPGEM Visual Results

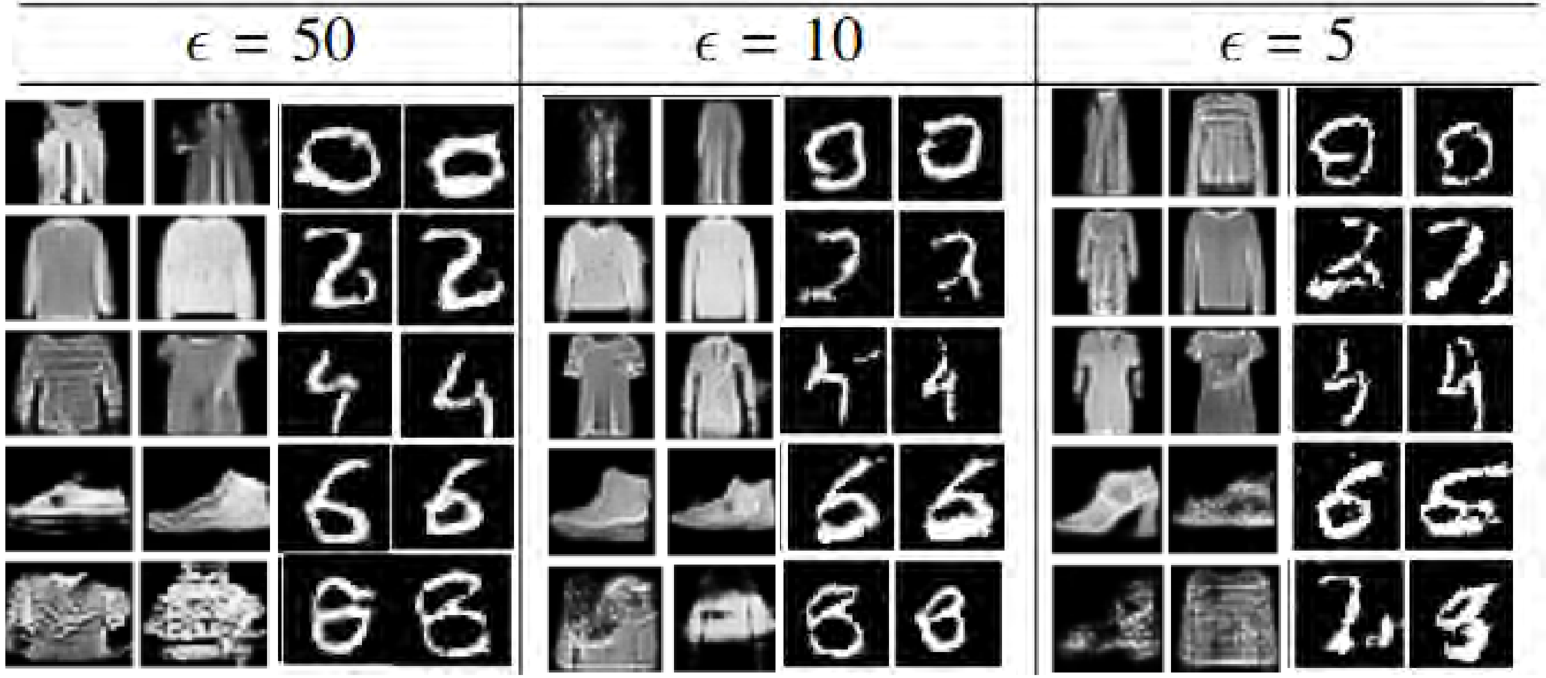

MNIST:

Public: 1, 3, 5, 7, 9

Private: 0, 2, 4, 6, 8

Fashion-MINIST

Public: T-shirt, Pullover, Dress, Sandal, Ankle Boot

Private: Trouser, Coast, Shirt, Sneaker, Bag

Conclusions

- Big big data
  - Expert + machine-learned features
  - Multi-resolution approaches
- Costly Annotation
  - Alternative to supervised learning: Semi-supervised learning and Active learning
  - Adaptation to class imbalance and other real-world problems
- Privacy
  - Local Synthetic Model: cleanest but some impact on downstream performance
  - Latent space processing
    - Latent-space DP GAN with model inversion
    - Exponential mechanism to sample latent vectors
- Applications:
  - early ASD risk based on behavior markers in videos,
  - WSI segmentation of brain tissues

Questions?