BTAS 2019

Session III: Unknown Attacks, Additional Sensors and Practical Tips

Host: Xiaoming Liu

MICHIGAN STATE UNIVERSITY

Computer Vision Lab

Outline

- Training-Testing Difference
- Alternative/Additional Sensors
- Practical Tips
- Future

Outline

- Training-Testing Difference
- Alternative/Additional Sensors
- Practical Tips
- Future

Training-Testing Difference

BTAS 2019

The testing scenarios are different with the training phase.

- Environment (Lighting, Indoor/outdoor, etc.)
- Camera/Image quality
- Subjects (Age, Race, etc.)
- Spoof types

Training-Testing Difference

BTAS 2019

The testing scenarios are different with the training phase.

- Environment (Lighting, Indoor/outdoor, etc.)
- Camera/Image quality
- Subjects (Age, Race, etc.)
- Spoof types

Cross-database Domain Adaption

Cross-database Domain Adaption

Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Multi-adversarial Discriminative Deep Domain Generalization, CVPR, 2019

Domain Adaptation in Multi-Channel Autoencoder based Features for Robust Face Anti-Spoofing, ICB 2019

Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation, ICB 2019

Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing

BTAS 2019

- Learn face anti-spoofing and face recognition at the same time
- Apply a Fast Domain Adaption (FDA) to remove the bias of different domain
- Share the weights of face anti-spoofing and face recognition

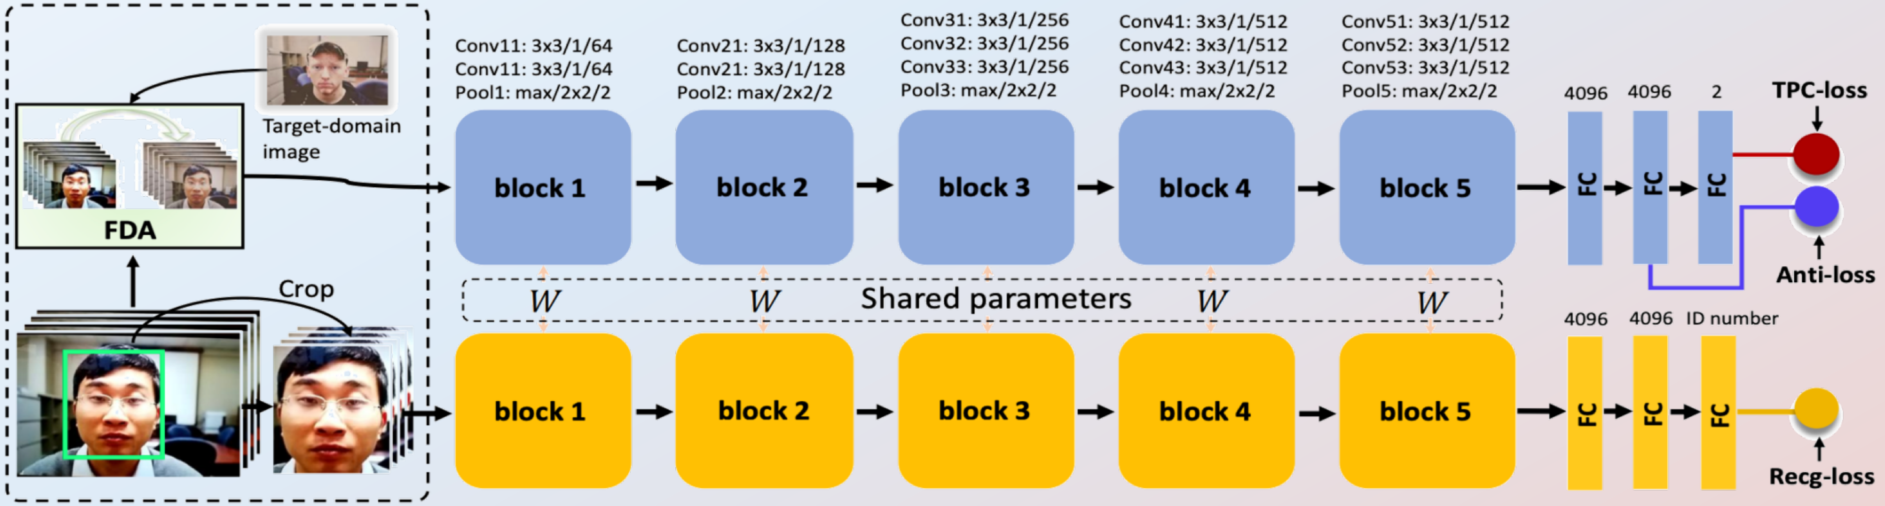

Li et. al., Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing

BTAS 2019

- Total Pairwise Confusion (TPC) loss

$$\mathcal{L}_{tpc}(\mathbf{x}_i, \mathbf{x}_j) = \sum_{i \neq j}^{M} ||\psi(\mathbf{x}_i) - \psi(\mathbf{x}_j)||_2^2$$

ψ(x) is the second fully connected layer of the face anti-spoofing branch

- Anti-loss: cross entropy losses for face anti-spoofing
- Recognition loss: cross entropy losses for face recognition

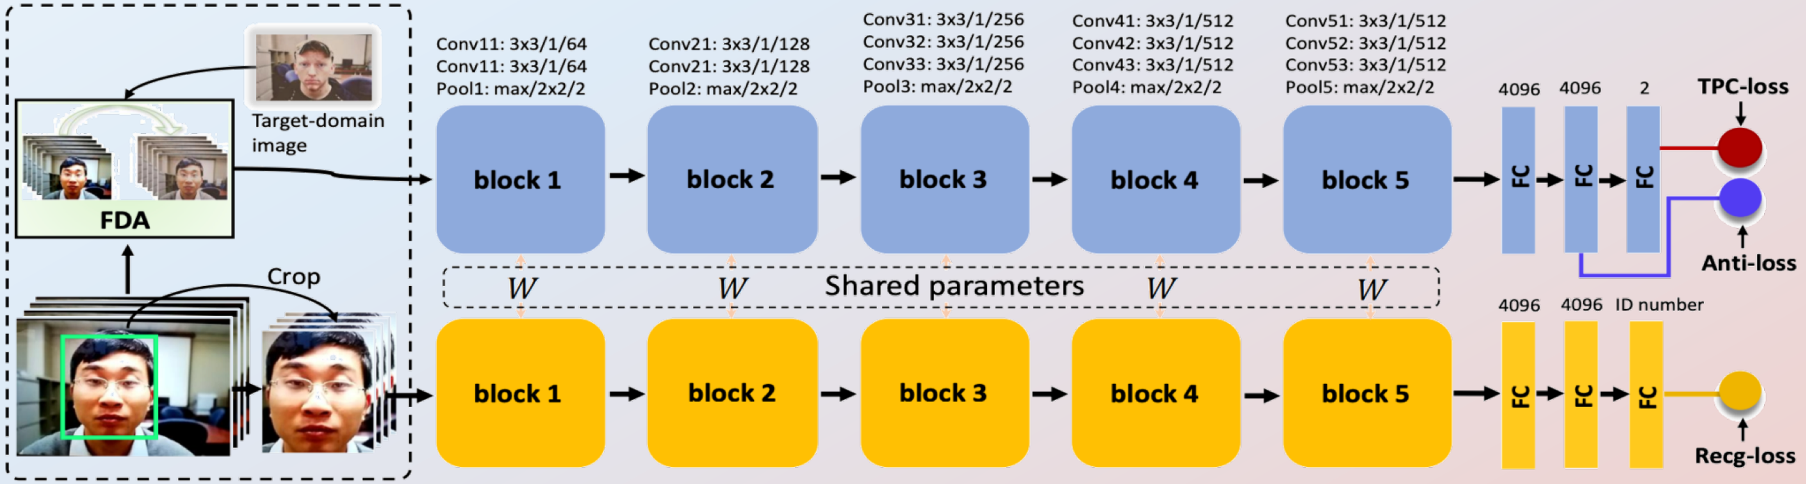

Li et. al., Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Feature w/ and w/o TPC loss

BTAS 2019

- Remove person id information from anti-spoofing feature
  - Irrelevant to face anti-spoofing
  - May lead to a more generalized feature

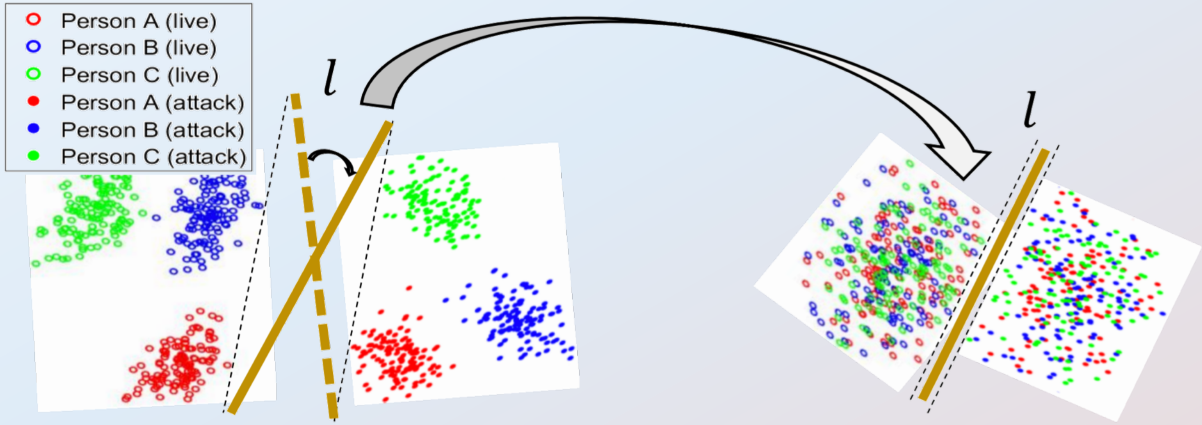

Feature distribution w/o $\mathcal{L}_{\mathrm{tpc}}$

Feature distribution w/ $\mathcal{L}_{\mathrm{tpc}}$

Li et. al., Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Feature w/ and w/o TPC loss

- Remove person id information from anti-spoofing feature
  - Irrelevant to face anti-spoofing
  - May lead to a more generalized feature

TPC/FDA	Intra-Test		Cross-Test	
	MFSD	Replay	MFSD $\rightarrow$ Replay	Replay $\rightarrow$ MFSD
$-$ $-$	10.5	0.6	39.4	34.6
$-$ $+$	11.2	0.6	36.3	38.3
$+$ $-$	**6.4**	**0**	28.5	26.6
$+$ $+$	8.3	0.3	**25.8**	**23.5**

Li et. al., Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing

- Fast Domain Adaption (FDA)
  - Style transfer network
  - Content loss + Style (domain) loss

$$\mathcal{L}_{\text{content}} = \frac{1}{C_j H_j W_j} ||\varphi_j(y) - \varphi_j(x)||_2^2$$

$$\mathcal{L}_{\text{domain}} = \frac{1}{C_j H_j W_j} ||G_j(y) - G_j(y_d)||_F^2$$

$$\hat{y} = \arg\min_P (\lambda_c \mathcal{L}_{\text{content}}(y, x) + \lambda_s \mathcal{L}_{\text{domain}}(y, y_d))$$

Live samples

Spoofing samples

Li et. al., Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Testing on Oulu

BTAS 2019

Protocol	Method	APCER	BPCER	ACER
P1	GRADIANT	1.3%	12.5%	6.9%
	Auxiliary	1.6%	**1.6%**	1.6%
	DS Net	**1.2%**	1.7%	**1.5%**
	GFA-CNN	2.5%	8.9%	5.7%
P2	Auxiliary	2.7%	2.7%	2.7%
	GRADIANT	3.1%	1.9%	2.5%
	DS Net	4.2%	4.4%	4.3%
	GFA-CNN	**2.5%**	**1.3%**	**1.9%**
P3	GRADIANT	**2.6±3.9%**	5.0±5.3%	3.8±2.4%
	Auxiliary	2.7±1.3%	**3.1±1.7%**	**2.9±1.5%**
	DS Net	4.0±1.8%	3.8±1.2%	3.6±1.6%
	GFA-CNN	4.3%	7.1%	5.7%
P4	GRADIANT	**5.0±4.5**%	15.0±7.1%	10.0±5.0%
	Auxiliary	9.3±5.6%	10.4±6.0%	9.5±6.0%
	DS Net	5.1±6.3%	**6.1±5.0%**	**5.6±5.7%**
	GFA-CNN	7.4%	10.4%	8.9%

Li et. al., Learning Generalizable and Identity-Discriminative Representations for Face Anti-Spoofing, TIFS, 2018

Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation

- Pretrain a source encoder/decoder
- Learn a target encoder such that discriminator cannot correctly predict the domain
- Classify with k-NN classifier

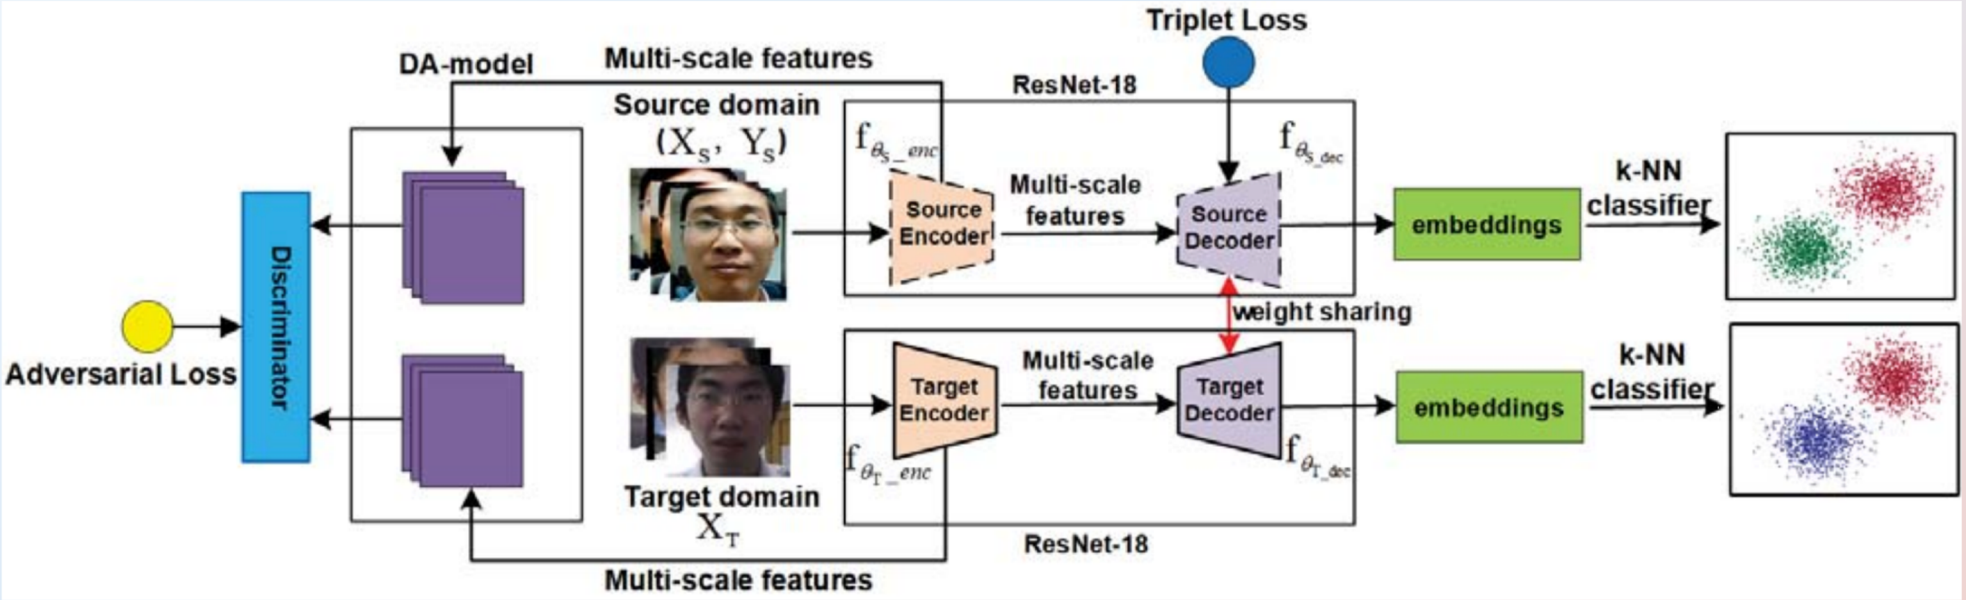

Wang et. al., Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation, ICB, 2019

Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation

- Encoder:
  - 4 convolution blocks
  - 1 pooling layer

- Decoder:
  - 2 fully connected layers

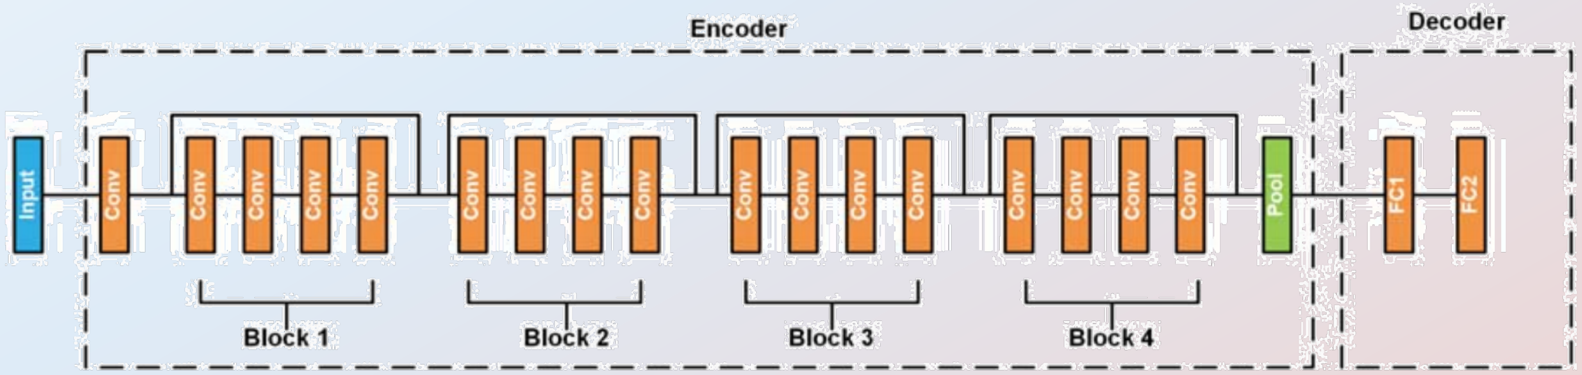

Wang et. al., Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation, 2019

Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation

BTAS 2019

Method	C → I	C → M	I → C	I → M	M → C	M → I	Average
Proposed w/o ML&ADA	43.8	33.8	49.5	41.3	45.4	39.6	42.2
Proposed w/o ML	43.7	29.6	50.0	35.4	46.5	38.7	40.7
Proposed w/o ADA	43.3	14.0	45.4	35.3	37.8	11.5	31.2
Proposed (full method)	**17.5**	**9.3**	**41.6**	**30.5**	**17.7**	**5.1**	**20.3**

Wang et. al., Improving Cross-database Face Presentation Attack Detection via Adversarial Domain Adaptation, 2019

Multi-adversarial Deep Domain Generalization for Face Presentation Attack Detection

- Learn a feature space that is discriminative and shared by multiple source domains

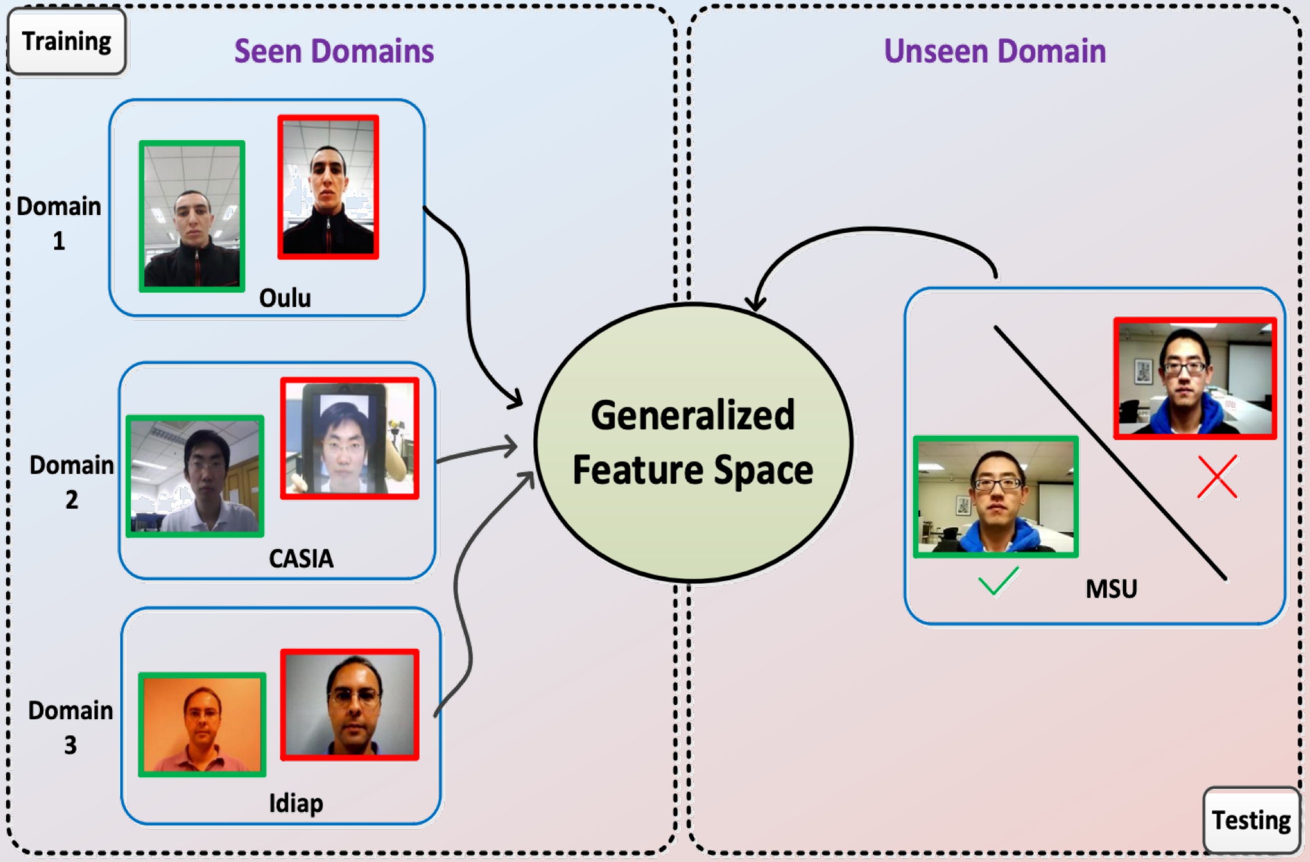

Shao et. al., Multi-adversarial Discriminative Deep Domain Generalization for Face Presentation Attack Detection, CVPR, 2019

Multi-adversarial Deep Domain Generalization for Face Presentation Attack Detection

BTAS 2019

- Feature generator
    - extract features for face anti-spoofing
    - adversarial-trained to remove domain information
- Depth estimation
    - improve the discriminativeness
- Dual-force triplet mining
    - enforce a smaller intra-class distance
    - enforce a larger inter-class distance
    - cross domain

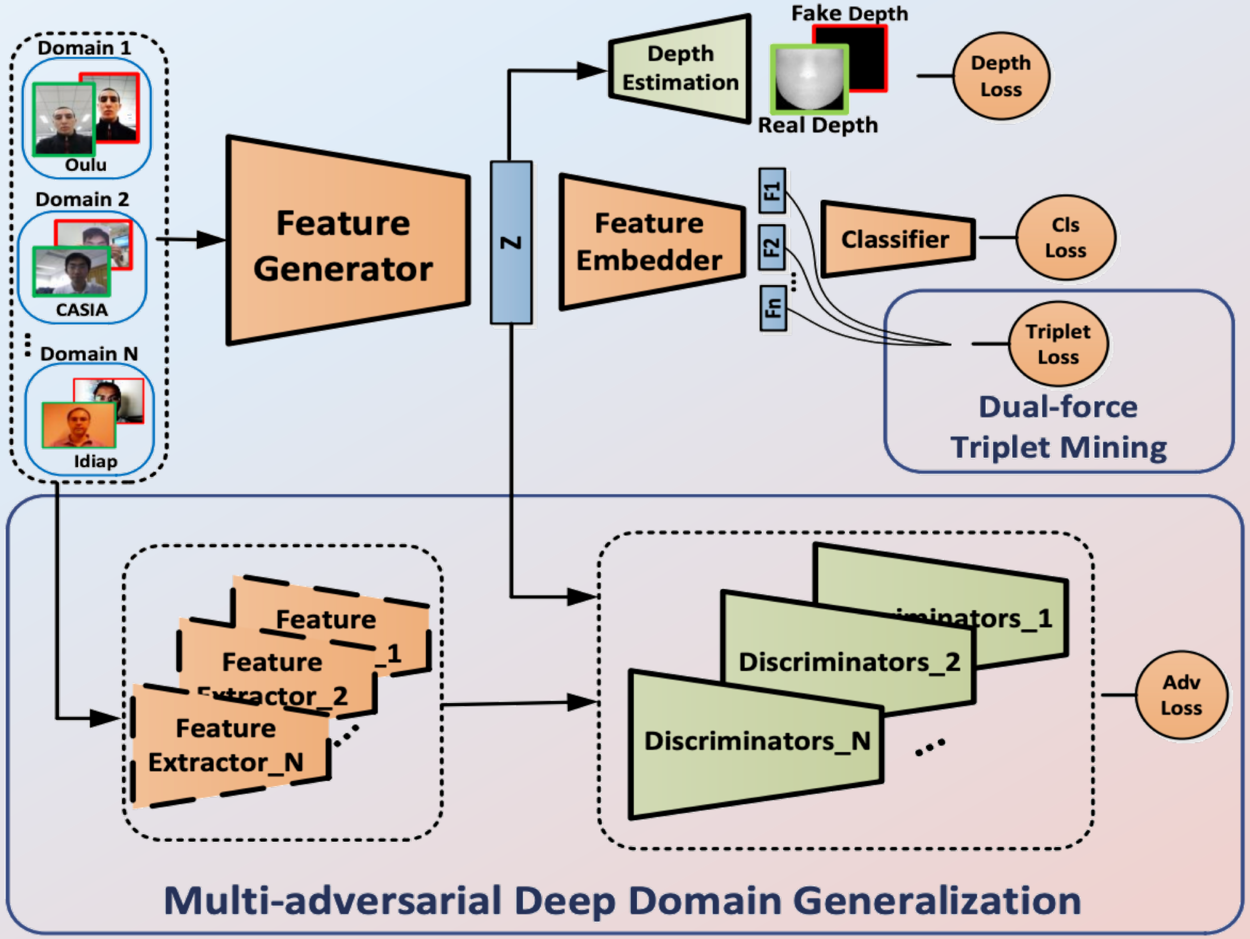

Shao et. al., Multi-adversarial Discriminative Deep Domain Generalization for Face Presentation Attack Detection, CVPR, 2019

Multi-adversarial Deep Domain Generalization for Face Presentation Attack Detection

BTAS 2019

- Learn features extractors for N domains
- Learn a feature generator for all domains
- Adversarial train N discriminators to make the feature generator more generalized.

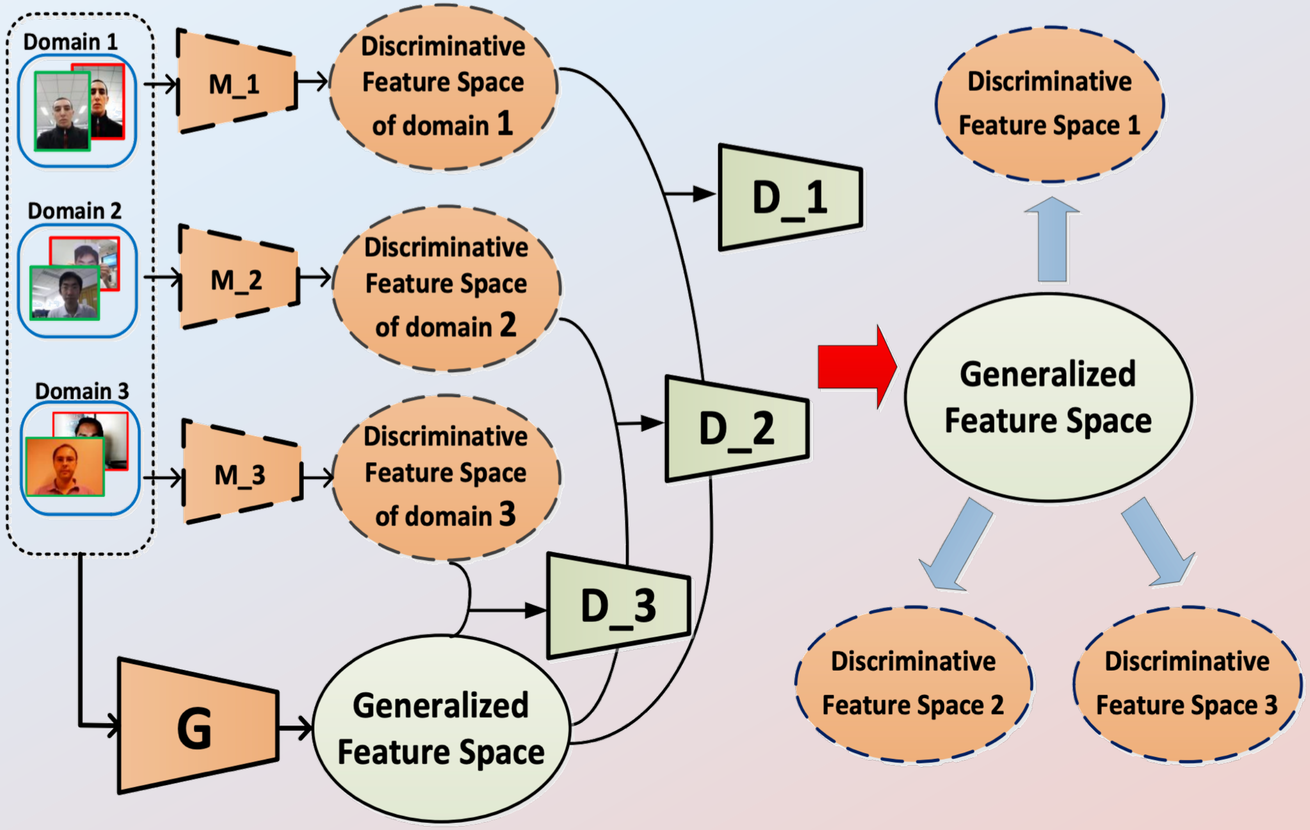

Shao et. al., Multi-adversarial Discriminative Deep Domain Generalization for Face Presentation Attack Detection, CVPR, 2019

Dual-force Triplet Mining

- In one domain
  - Minimize live-to-live / spoof-to-spoof distance between different subjects
  - Maximize live-to-spoof distance between different subjects
- Cross domains
  - Minimize live-to-live / spoof-to-spoof distance between different subjects
  - Maximize live-to-spoof distance between different subjects

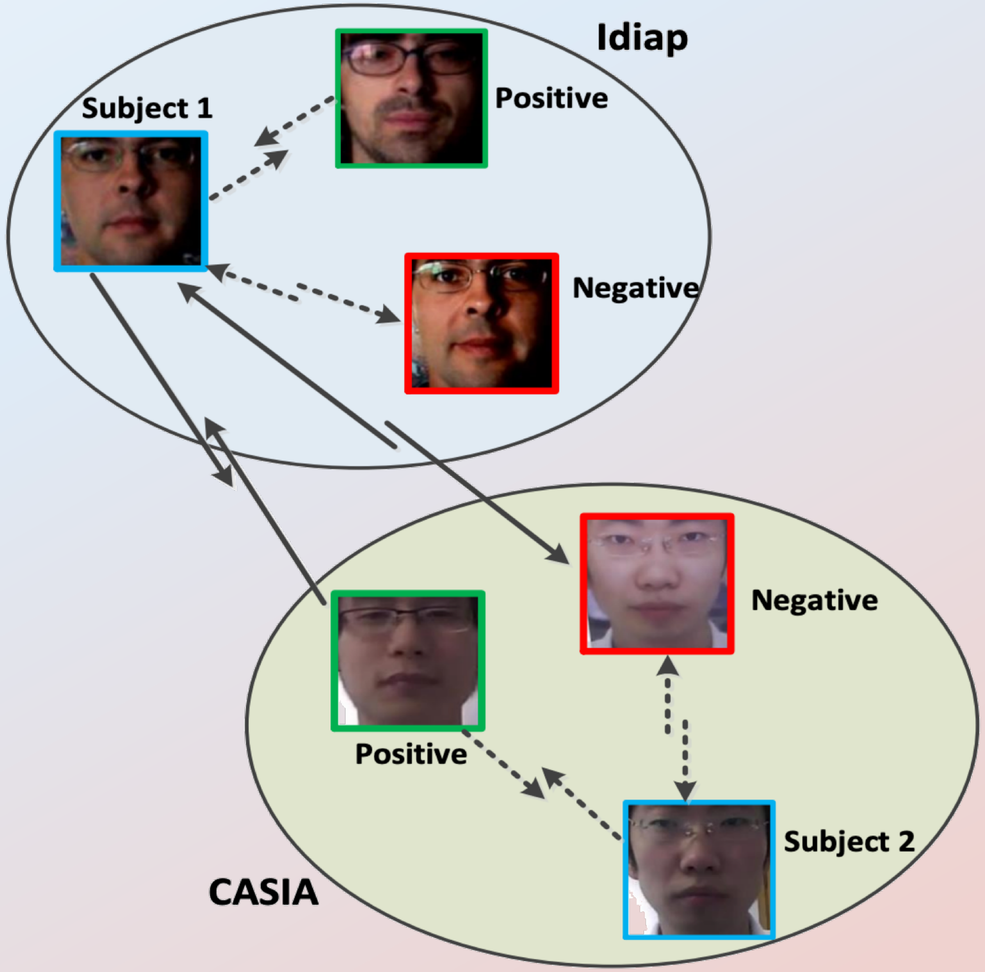

Shao et. al., Multi-adversarial Discriminative Deep Domain Generalization for Face Presentation Attack Detection, CVPR, 2019

Multi-adversarial Discriminative Deep Domain Generalization

BTAS 2019

Method	O&C&I to M		O&M&I to C		O&C&M to I		I&C&M to O	
	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)	AUC(%)
MS_LBP	29.76	78.50	54.28	44.98	50.30	51.64	50.29	49.31
Binary CNN	29.25	82.87	34.88	71.94	34.47	65.88	29.61	77.54
IDA	66.67	27.86	55.17	39.05	28.35	78.25	54.20	44.59
Color Texture	28.09	78.47	30.58	76.89	40.40	62.78	63.59	32.71
LBPTOP	36.90	70.80	42.60	61.05	49.45	49.54	53.15	44.09
Auxiliary(Depth Only)	22.72	85.88	33.52	73.15	29.14	71.69	30.17	77.61
Auxiliary(All)	–	–	28.4	–	27.6	–	–	–
Ours (MADDG)	**17.69**	**88.06**	**24.5**	**84.51**	**22.19**	**84.99**	**27.98**	**80.02**

Method	O&C&I to M		O&M&I to C		O&C&M to I		I&C&M to O	
	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)	AUC(%)
MMD-AAE	27.08	83.19	44.59	58.29	31.58	75.18	40.98	63.08
Ours (MADDG)	**17.69**	**88.06**	**24.5**	**84.51**	**22.19**	**84.99**	**27.98**	**80.02**

Shao et. al., Multi-adversarial Discriminative Deep Domain Generalization, CVPR, 2019

Domain Adaptation in Multi-Channel Autoencoder based Features for Robust Face Anti-Spoofing

BTAS 2019

- Use multi-modality data (RGB, NIR, and Depth) instead of RGB only
- Domain Adaption: fine-tuning (RGB → NIR-Depth)

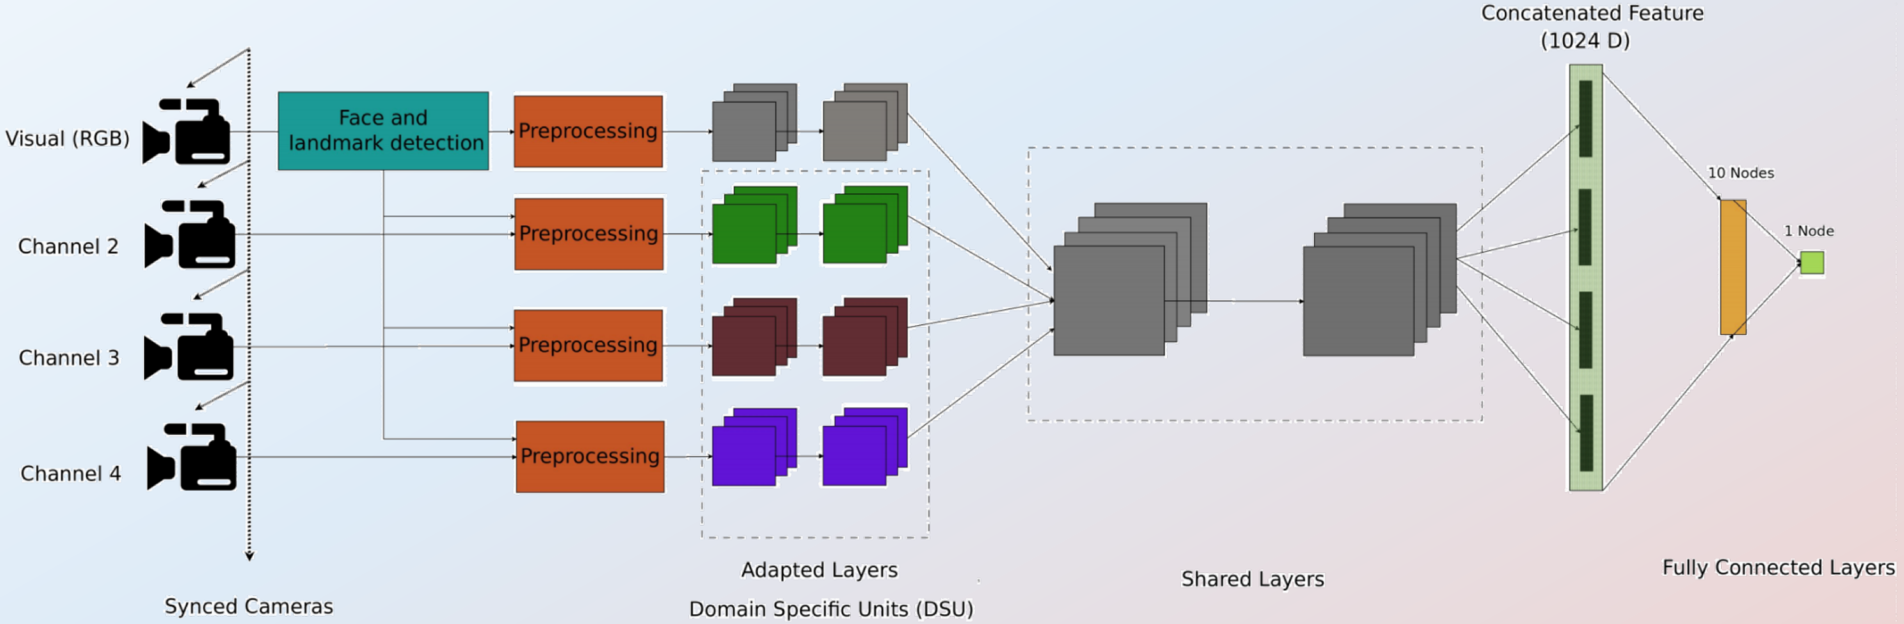

George et. al., Biometric Face Presentation Attack Detection with Multi-Channel Convolutional Neural Network, TIFS 2019

Domain Adaptation in Multi-Channel Autoencoder based Features for Robust Face Anti-Spoofing

BTAS 2019

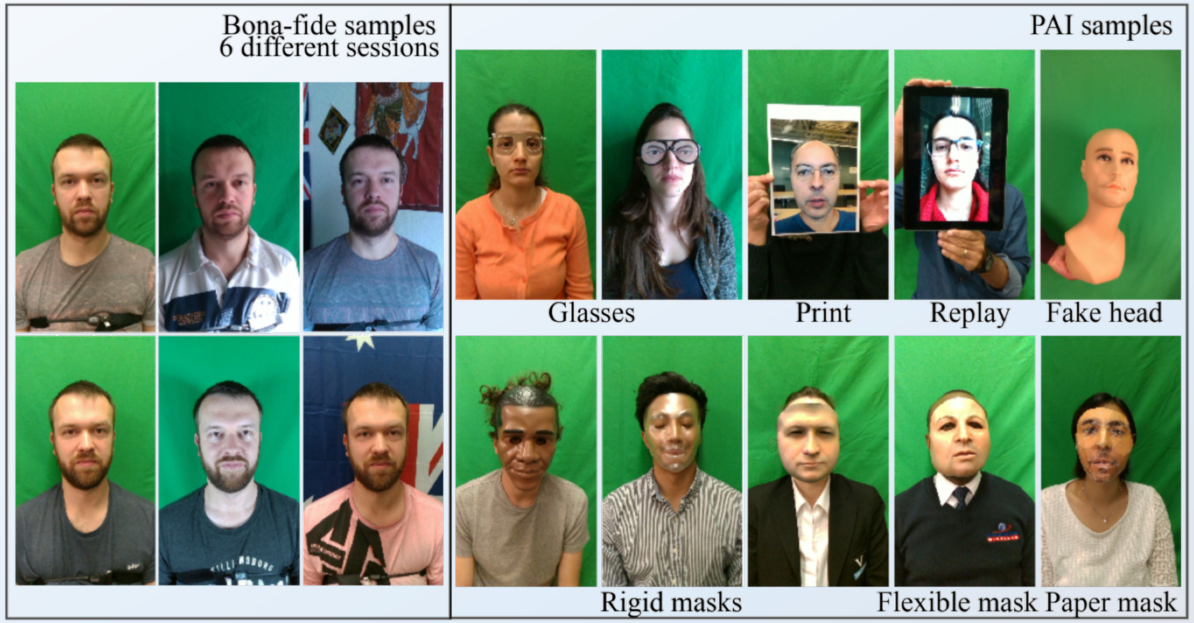

Method	dev (%)		test (%)		
	APCER	ACER	APCER	BPCER	ACER
Color (IQM-LR)	76.58	38.79	87.49	0	43.74
Depth (LBP-LR)	57.71	29.35	65.45	0.03	32.74
Infrared (LBP-LR)	32.79	16.9	29.39	1.18	15.28
Thermal (LBP-LR)	11.79	6.4	16.43	0.5	8.47
Score fusion (IQM-LBP-LR Mean fusion)	10.52	5.76	13.92	1.17	7.54
Color (RDWT-Haralick-SVM)	36.02	18.51	35.34	1.67	18.5
Depth (RDWT-Haralick-SVM)	34.71	17.85	43.07	0.57	21.82
Infrared (RDWT-Haralick-SVM)	14.03	7.51	12.47	0.05	6.26
Thermal (RDWT-Haralick-SVM)	21.51	11.26	24.11	0.85	12.48
Score fusion (RDWT-Haralick-SVM Mean fusion)	6.2	3.6	6.39	0.49	3.44
FASNet	18.89	9.94	17.22	5.65	11.44

George et. al., Biometric Face Presentation Attack Detection with Multi-Channel Convolutional Neural Network, TIFS 2019

Unknown Attack Detection

- One-class SVM
- Gaussian Mixture Model
- AutoEncoder

Unknown Attack Detection

BTAS 2019

An Anomaly Detection Approach to Face Spoofing Detection: A New Formulation and Evaluation Protocol, IEEE Access, 2017

Unknown Presentation Attack Detection with Face RGB Images, ICB, 2018

Deep Anomaly Detection for Generalized Face Anti-Spoofing, CVPRW, 2019

Deep Tree Learning for Zero-shot Face Anti-Spoofing, CVPR 2019

An Anomaly Detection Approach to Face Spoofing Detection: A New Formulation and Evaluation Protocol

A very comprehensive study on various hand-crafted feature and classifiers.

- Feature: LBP-TOP, LPQ-TOP, BSIF-TOP, Image quality measures
- Classifier: SVM1, SVM2, LDA2, Sparse representation classifier (SRC)1, SRC 2
- Dataset: CASIA-FASD, Replay-attack, MSU-MFSD

Arashlool et. al., An Anomaly Detection Approach to Face Spoofing Detection: A New Formulation and Evaluation Protocol, 2017

An Anomaly Detection Approach to Face Spoofing Detection: A New Formulation and Evaluation Protocol

BTAS 2019

A very comprehensive study on various hand-crafted feature and classifiers.

- Feature: LBP-TOP, LPQ-TOP, BSIF-TOP, Image quality measures
- Classifier: SVM1, SVM2, LDA2, Sparse representation classifier (SRC)1, SRC 2
- Dataset: CASIA-FASD, Replay-attack, MSU-MFSD

- Conclusion: neither the two-class systems nor the one-class approaches perform well enough

Arashlool et. al., An Anomaly Detection Approach to Face Spoofing Detection: A New Formulation and Evaluation Protocol, 2017

Unknown Presentation Attack Detection with Face RGB Images

A very comprehensive study on various hand-crafted feature and classifiers.

- Feature: Color LBP
- Classifier: SVM1, Auto Encoder, GMM
- Dataset: CASIA-FASD, Replay-attack, MSU-MFSD

Xiong et. al., Unknown Presentation Attack Detection with Face RGB Images, ICB, 2018

Unknown Presentation Attack Detection with Face RGB Images

BTAS 2019

	CASIA			Replay-Attack			MSU			All	
	Video	Cut Photo	Warped Photo	Video	Digital Photo	Printed Photo	Printed Photo	HR Video	Mobile Video	Mean	Std
OC-SVM$_{RBF}$ + IMQ[1]	68.89	61.95	74.80	98.24	90.82	53.23	63.94	63.00	76.38	72.80	14.48
OC-SVM$_{RBF}$ + BSIF[1]	70.74	60.73	95.90	84.03	88.14	73.66	64.81	87.44	74.69	78.68	11.74
SVM$_{RBF}$ + LBP[5]	91.49	91.70	84.47	99.08	98.17	87.28	47.68	99.50	97.61	88.55	16.25
NN + LBP	94.16	88.39	79.85	99.75	95.17	78.86	50.57	99.93	93.54	86.69	15.56
GMM + LBP	90.91	77.52	62.61	93.20	87.80	89.19	68.18	91.21	94.04	83.85	11.60
OC-SVM$_{RBF}$ + LBP	91.21	82.32	65.58	91.55	84.97	87.19	71.46	96.89	93.57	84.97	10.42
AE + LBP	87.00	80.48	65.84	88.62	84.67	85.09	71.25	96.00	95.64	83.84	10.10

- Dataset: CASIA-FASD, Replay-attack, MSU-MFSD
- Conclusion: improve the performance
  - NN+LBP works best on C+R+M protocols
  - AE+LBP works best on Oulu protocols

Xiong et. al., Unknown Presentation Attack Detection with Face RGB Images, ICB, 2018

Deep Anomaly Detection for Generalized Face Anti-Spoofing

- Deep metric learning
- Triplet Focal loss
  - Focus on the harder cases

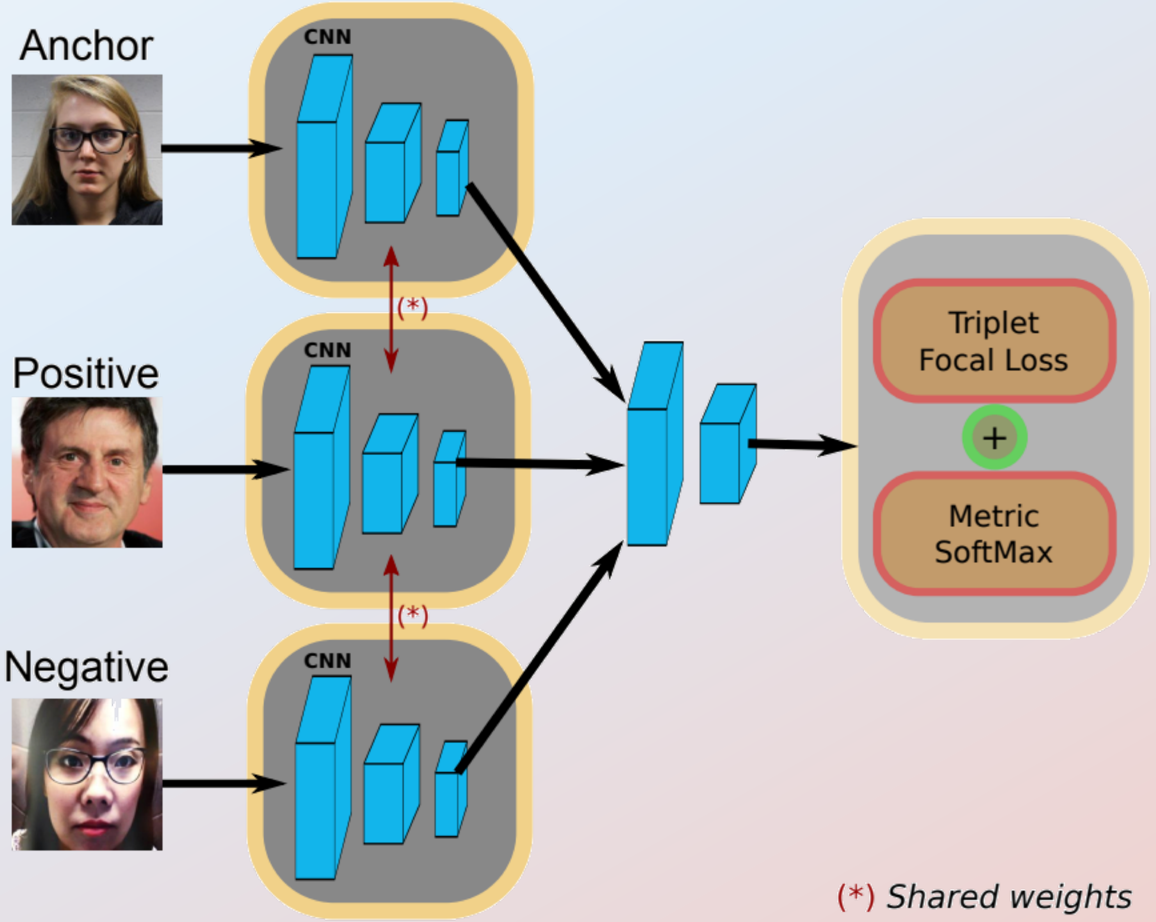

Perez-Cabo et. al., Deep Anomaly Detection for Generalized Face Anti-Spoofing, CVPRW, 2019

Literature and Issues

- Limited Spoof Types[1,2]
- Only model the live distribution[1,2]

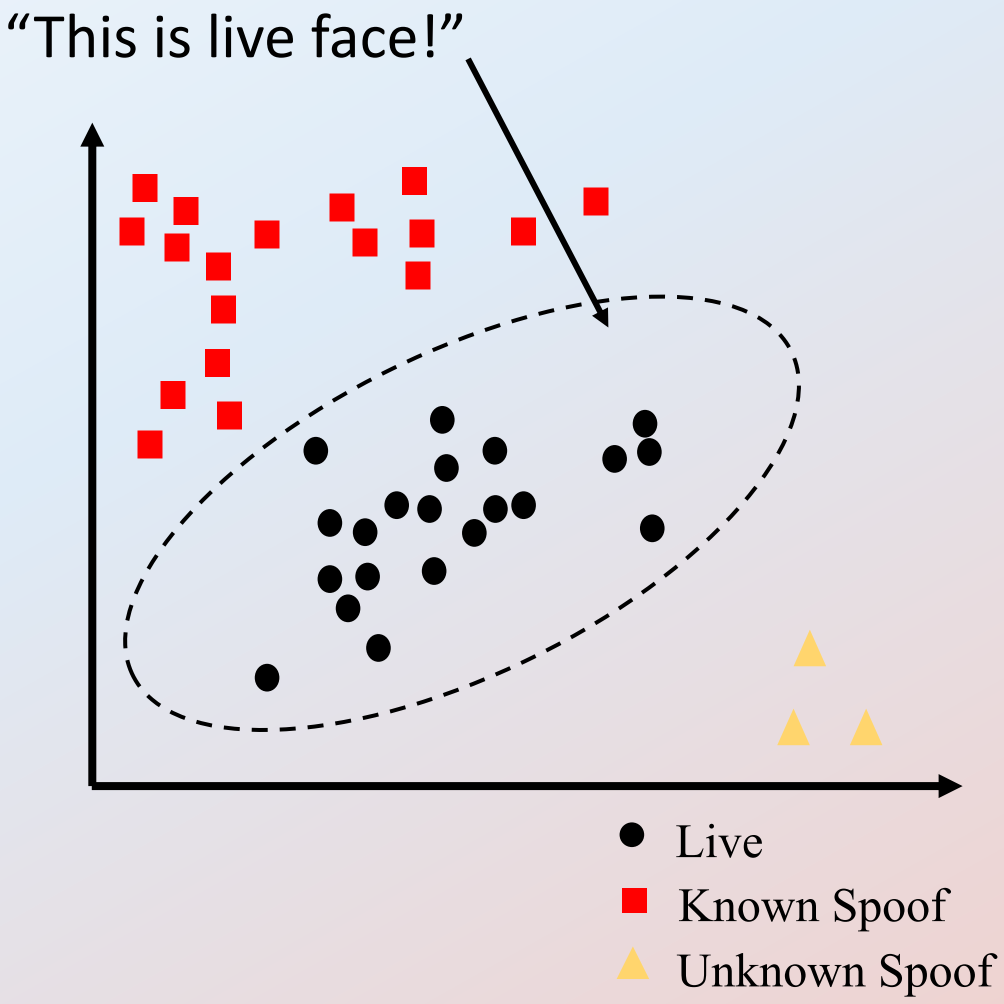

[1] S. R. Arashloo et. al. An anomaly detection approach to face spoofing detection: a new formulation and evaluation protocol.
[2] F. Xiong and W. Abdalmageed. Unknown presentation attack detection with face RGB images. BTAS 2018

What if More Spoof Types?

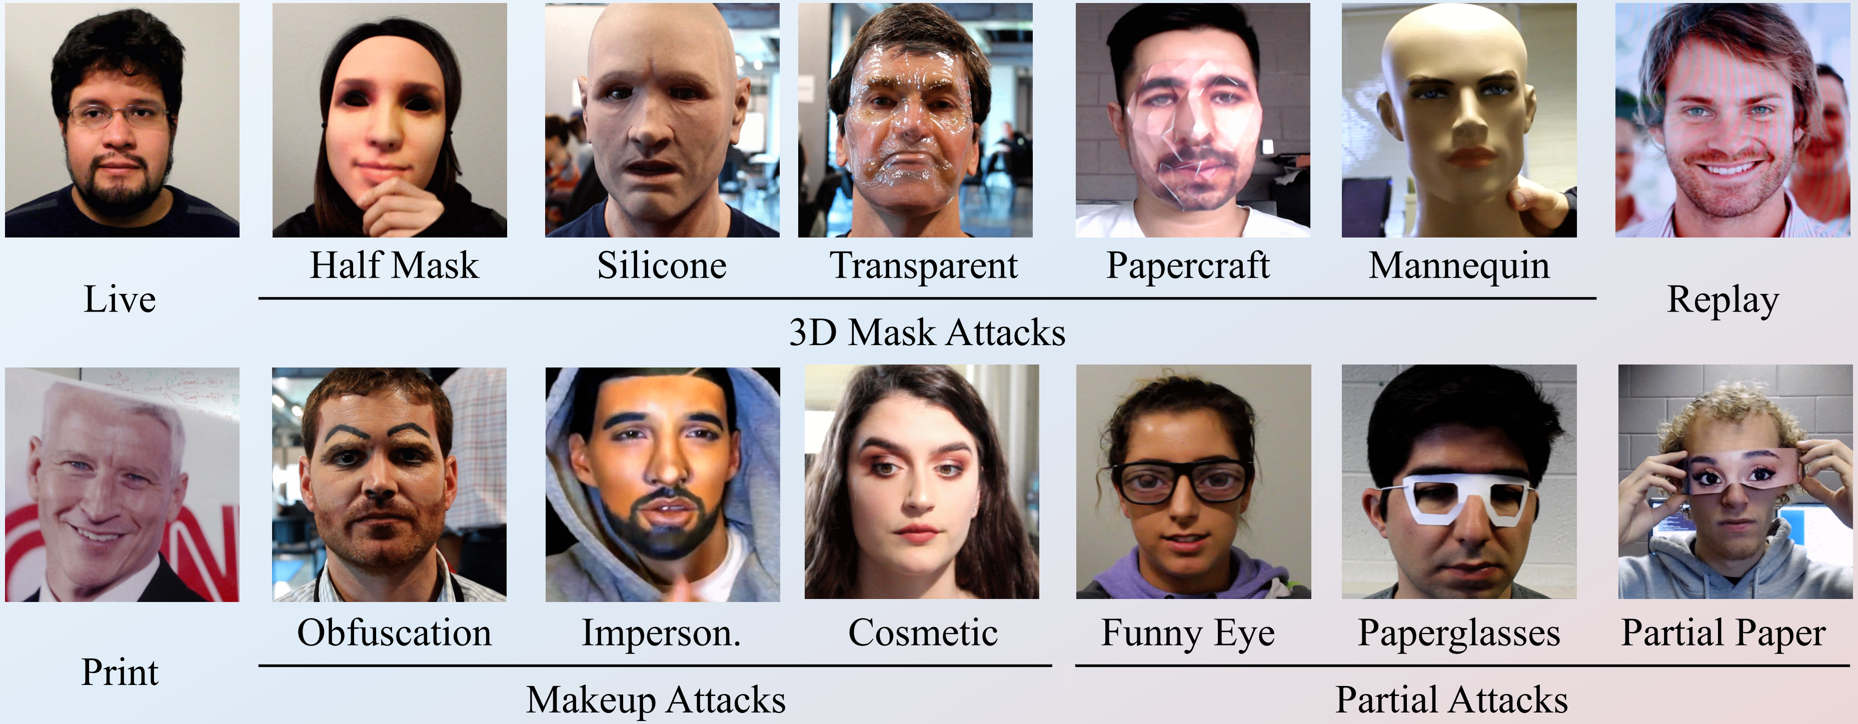

Deep Tree Learning for Zero-shot Face Anti-Spoofing

BTAS 2019

- Previous methods only model the live
- Learning semantic spoof attributes

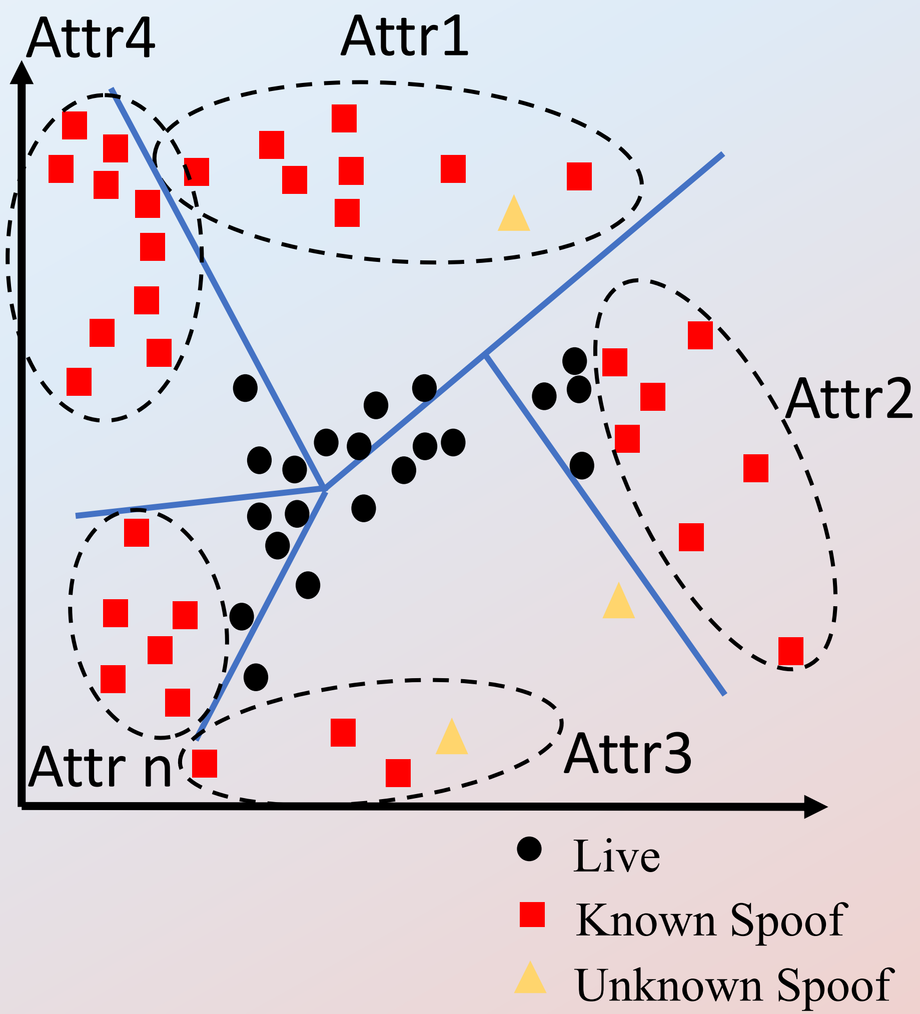

Liu et. al., Deep Tree Learning for Zero-shot Face Anti-Spoofing, CVPR 2019

Deep Tree Networks (DTN)

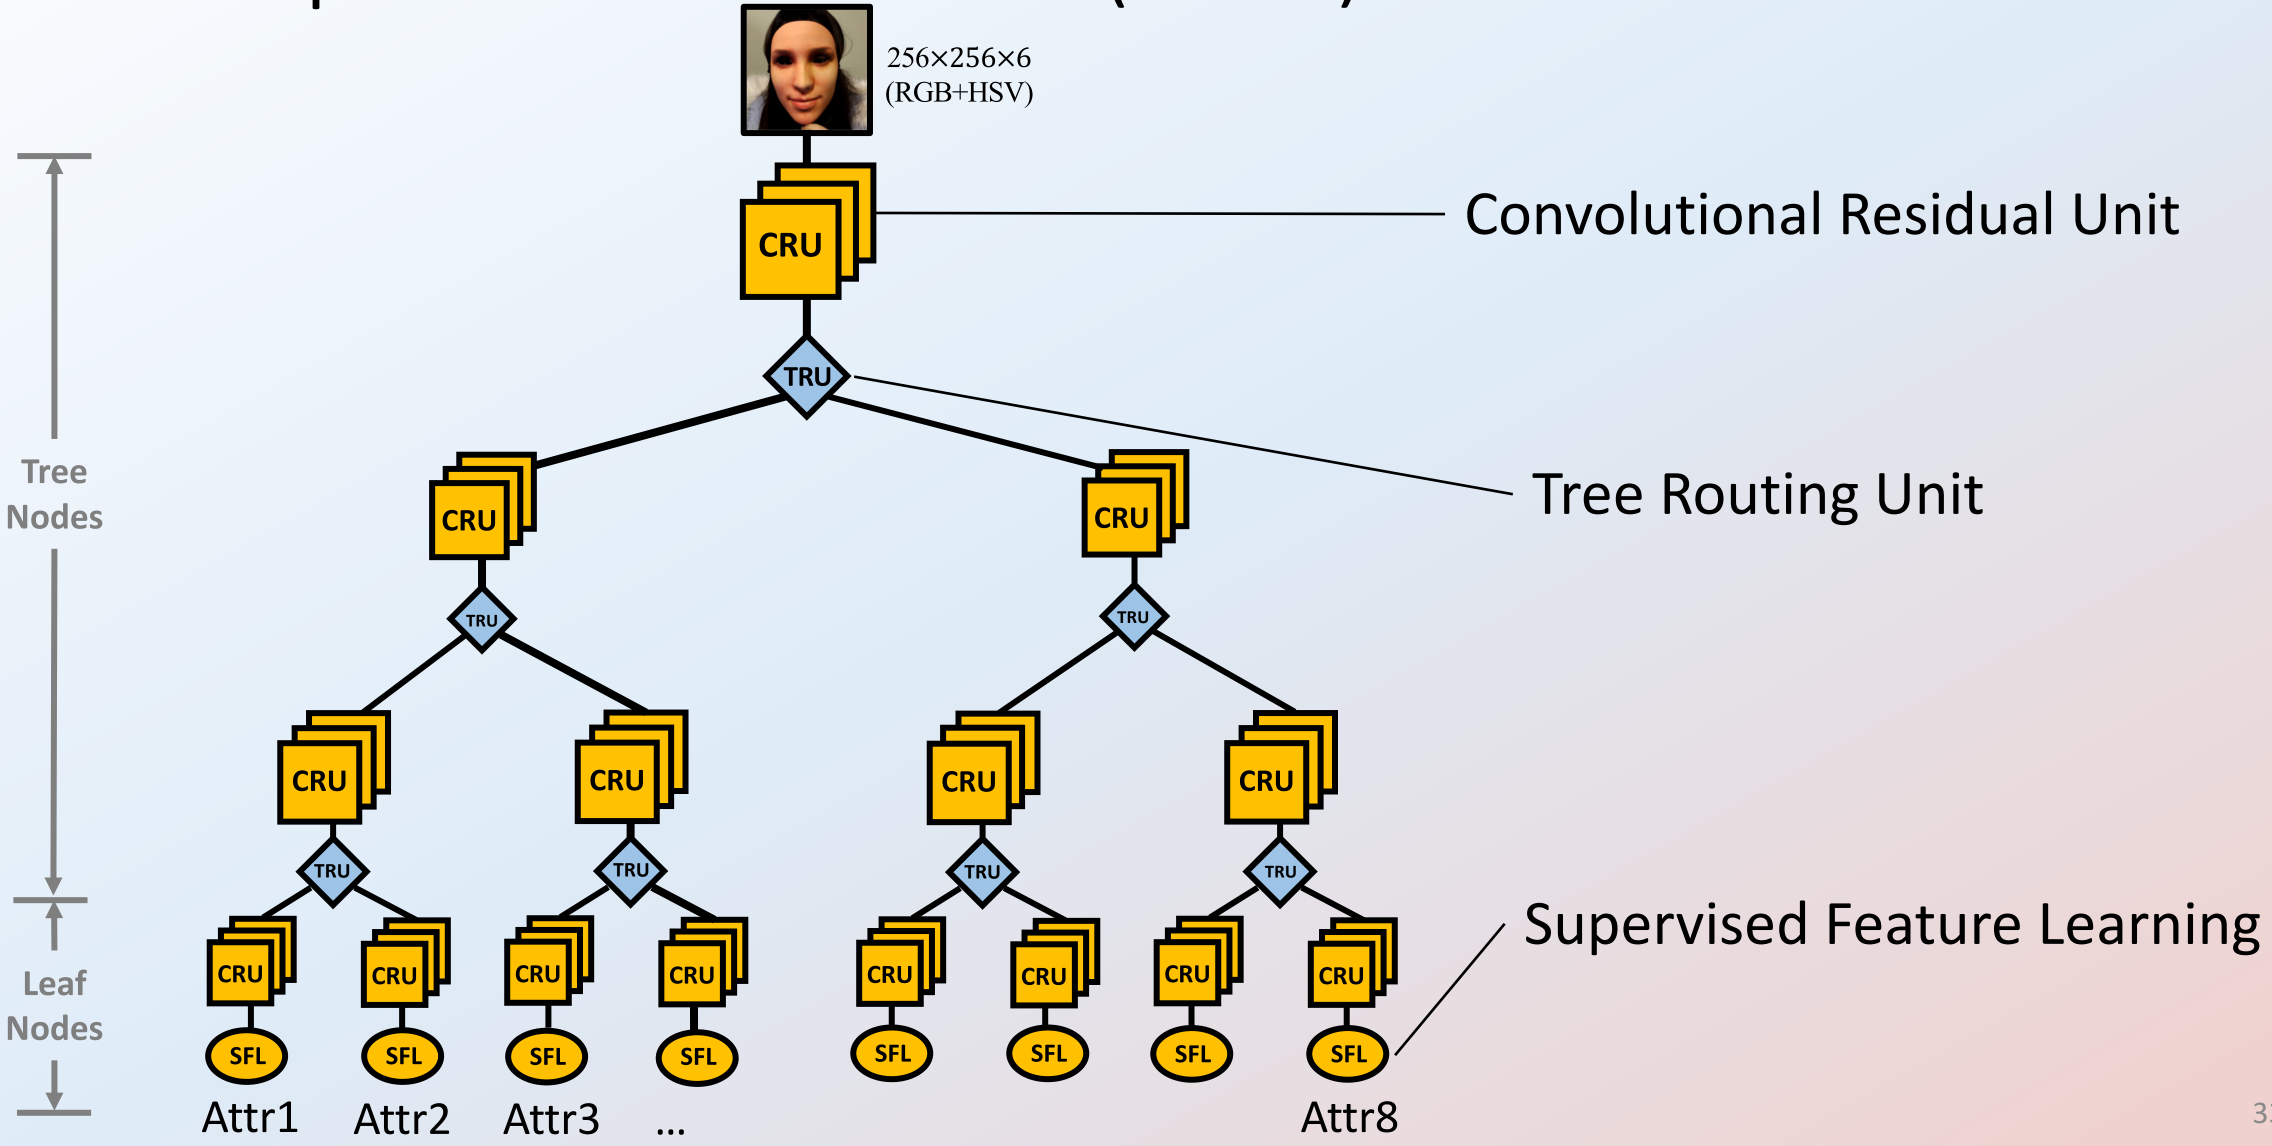

Deep Tree Networks (DTN)

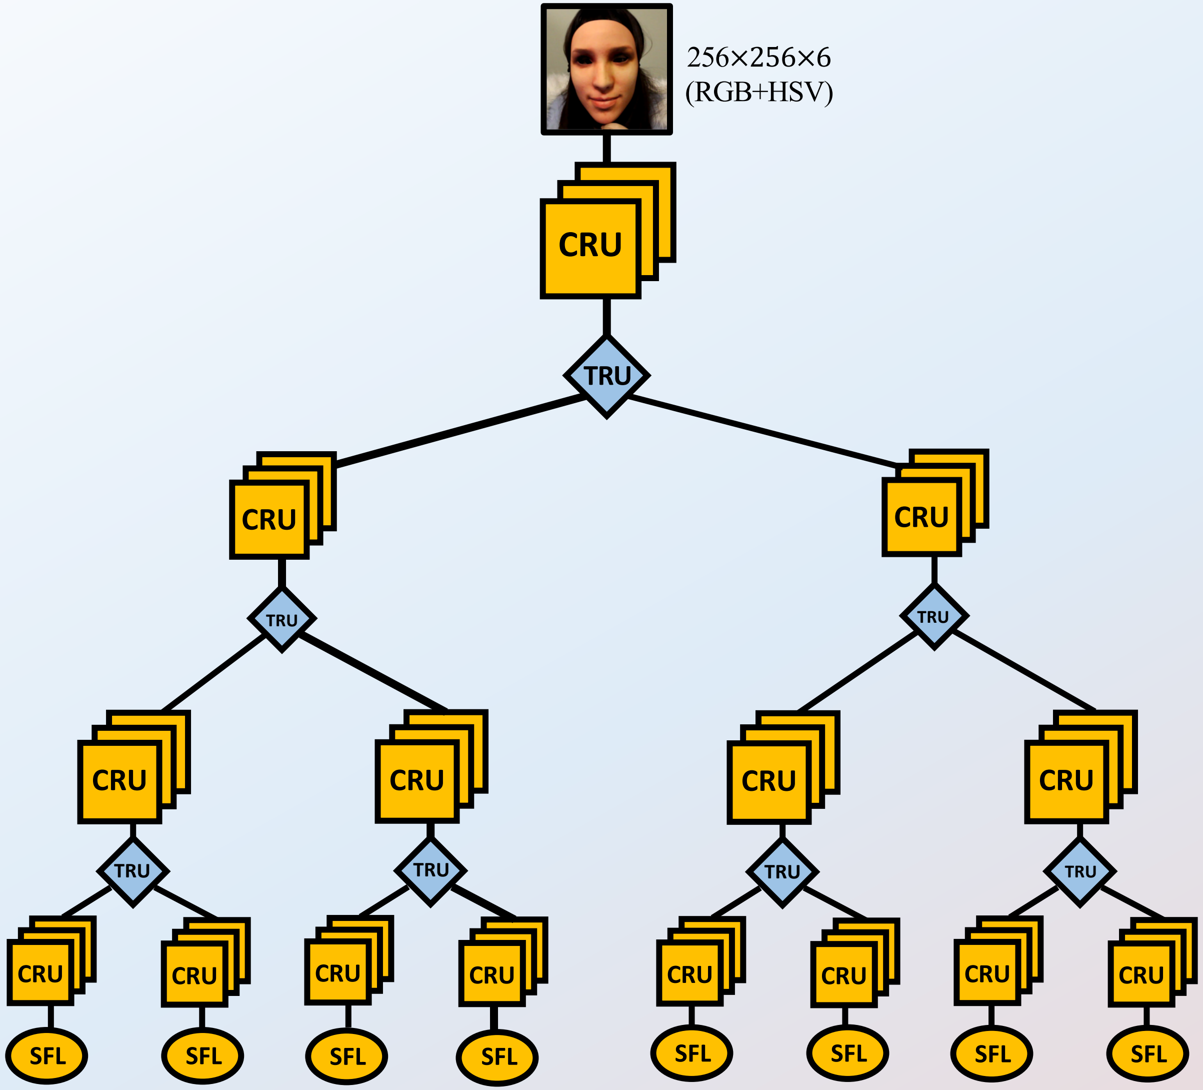

Deep Tree Networks (DTN)

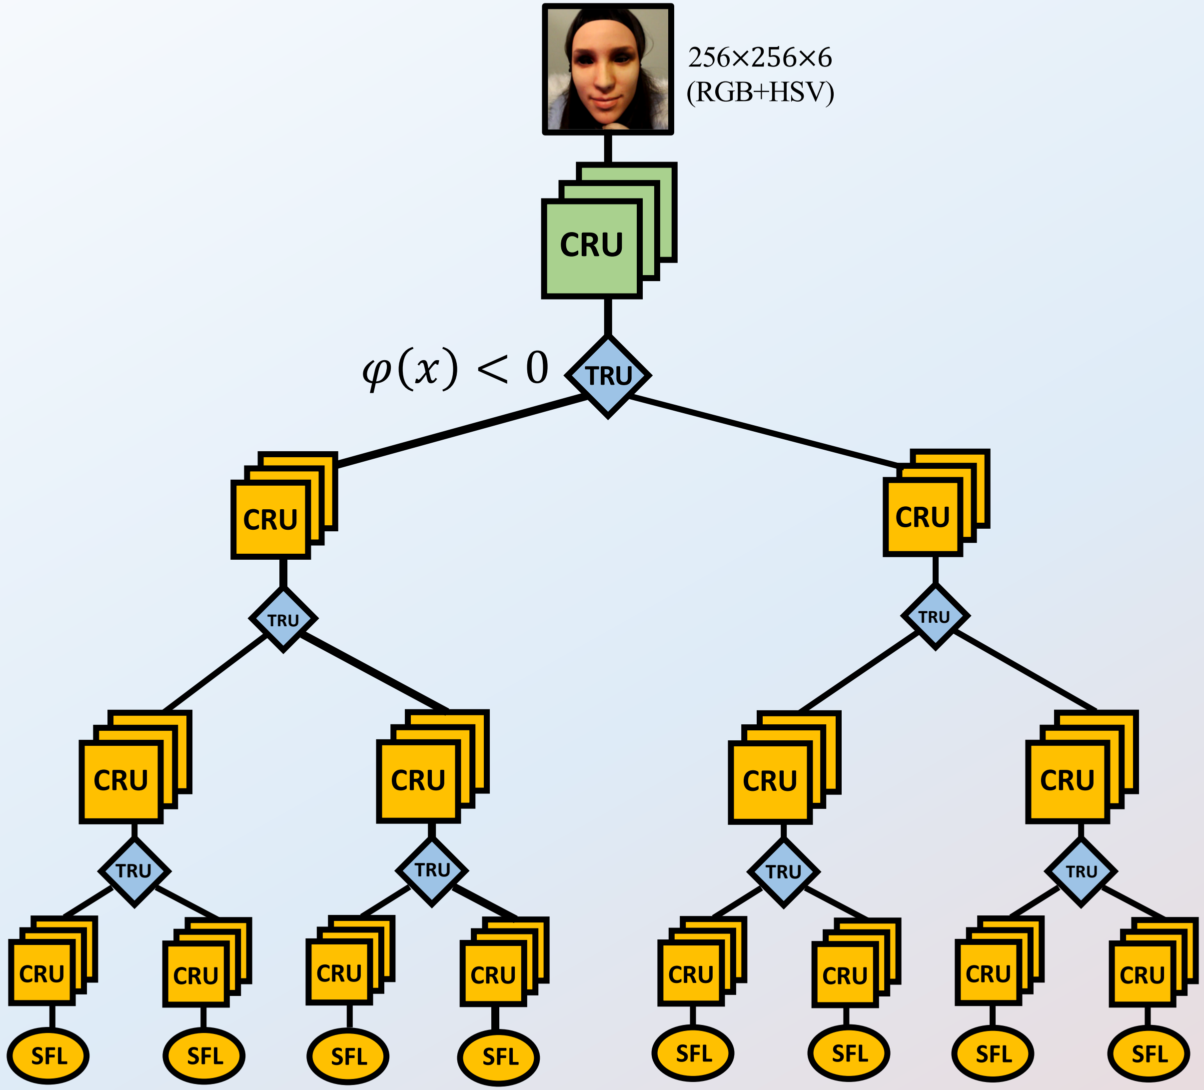

BTAS 2019

Deep Tree Networks (DTN)

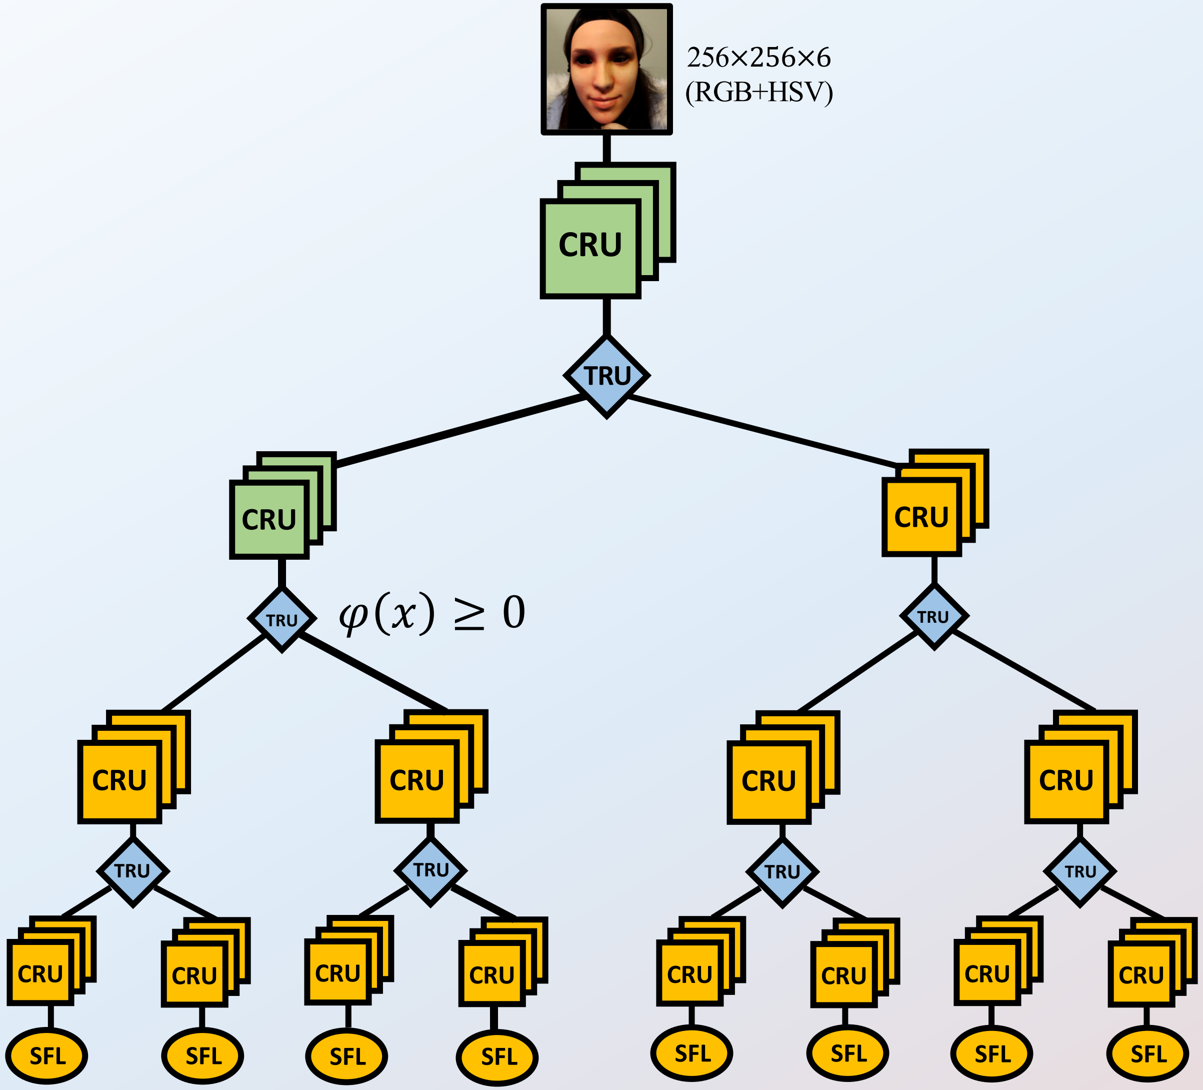

Deep Tree Networks (DTN)

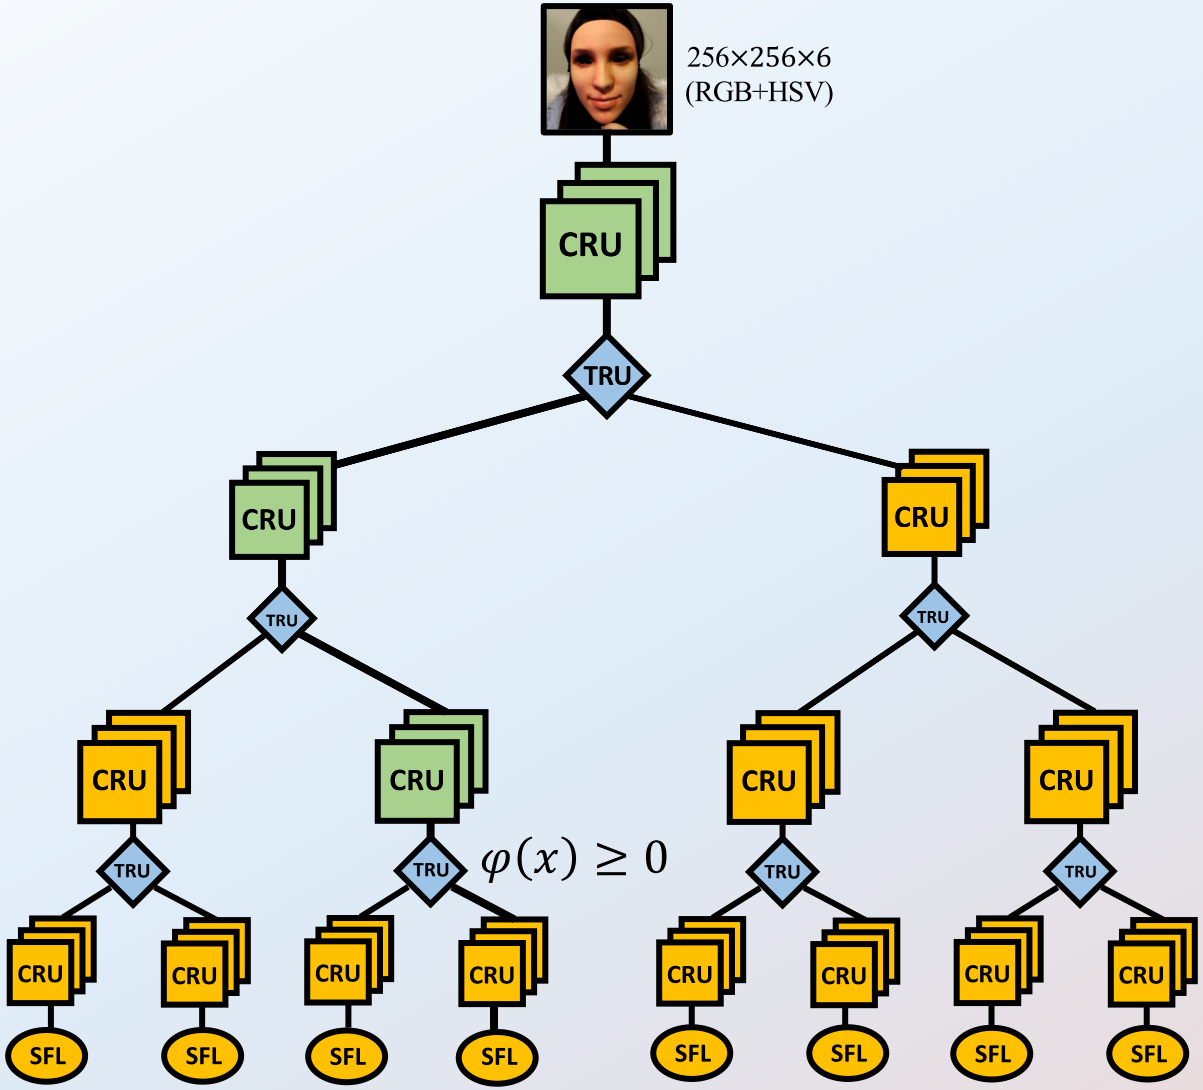

Deep Tree Networks (DTN)

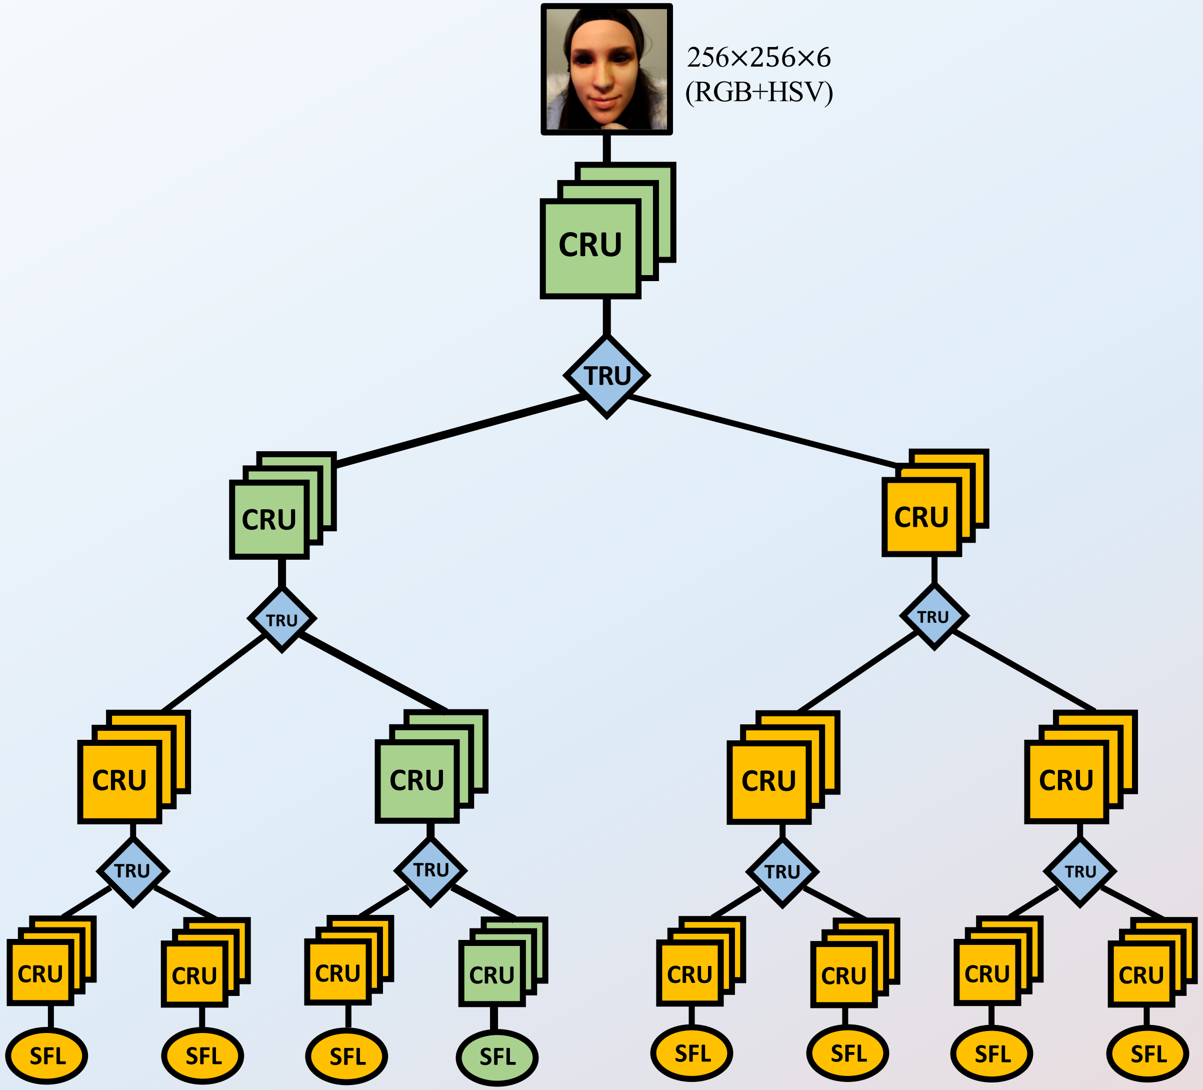

Supervised Feature Learning

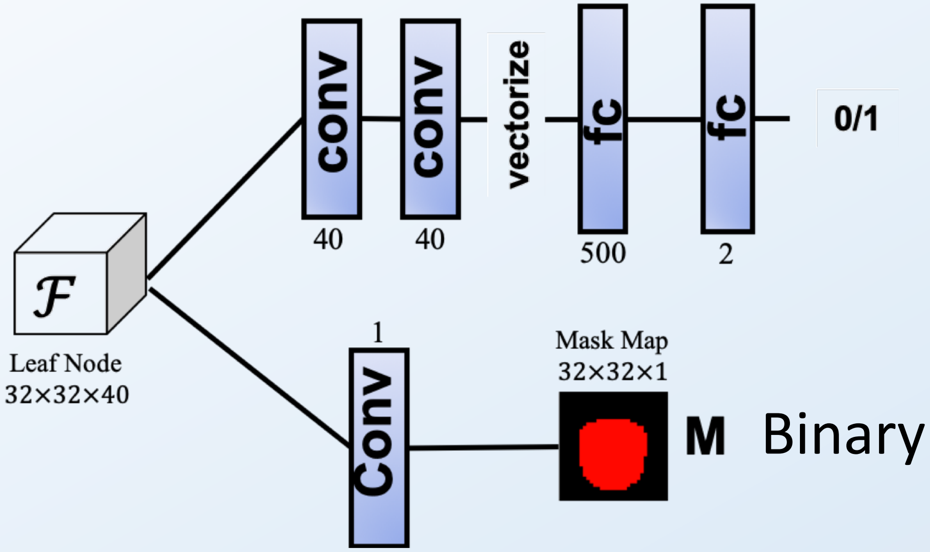

Classification

Binary Mask Regression

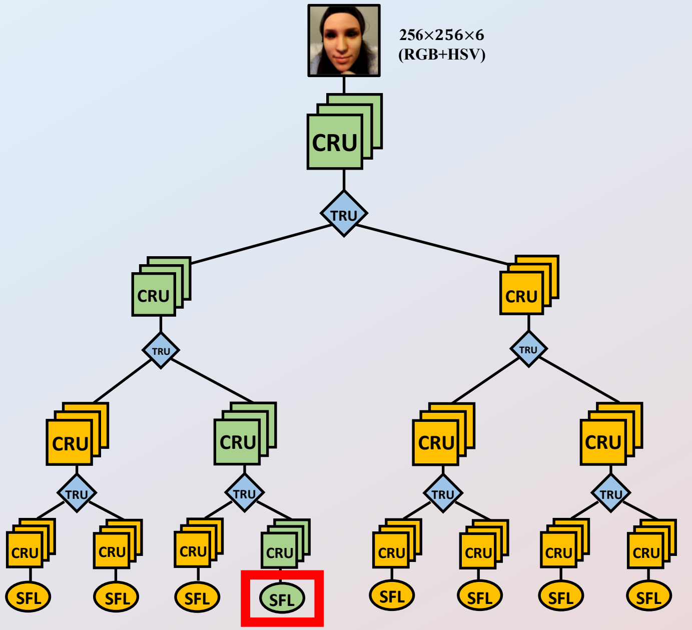

Supervised Feature Learning

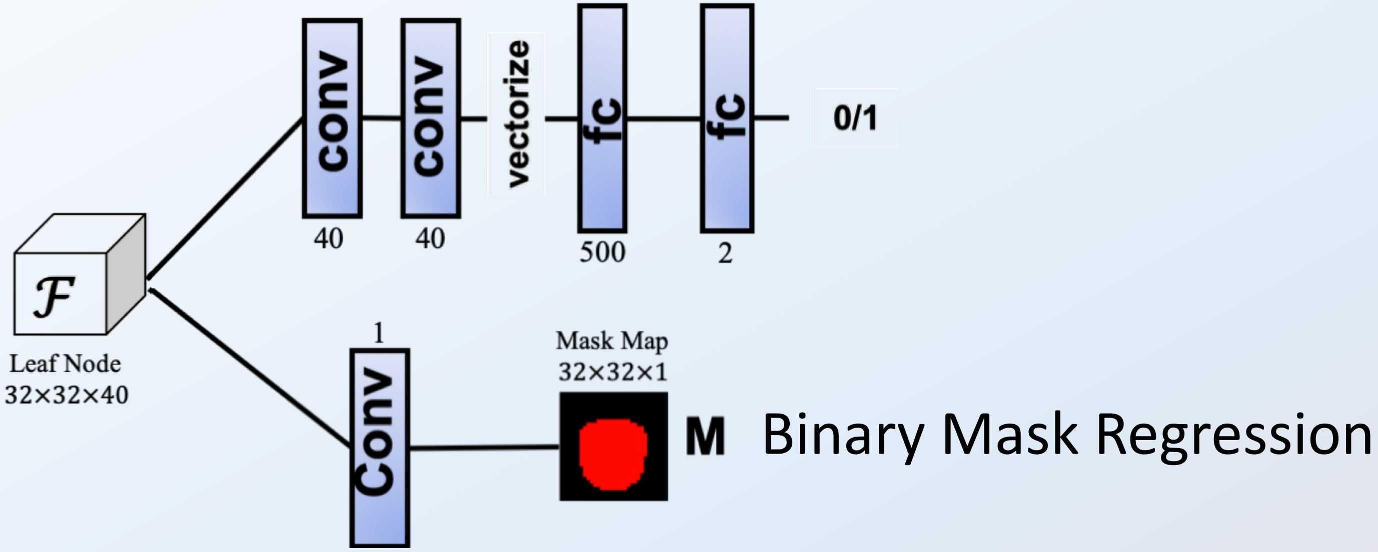

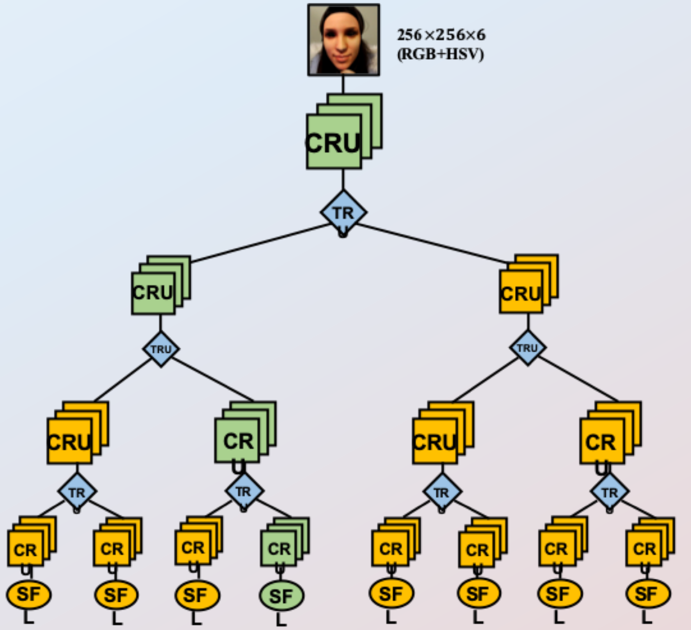

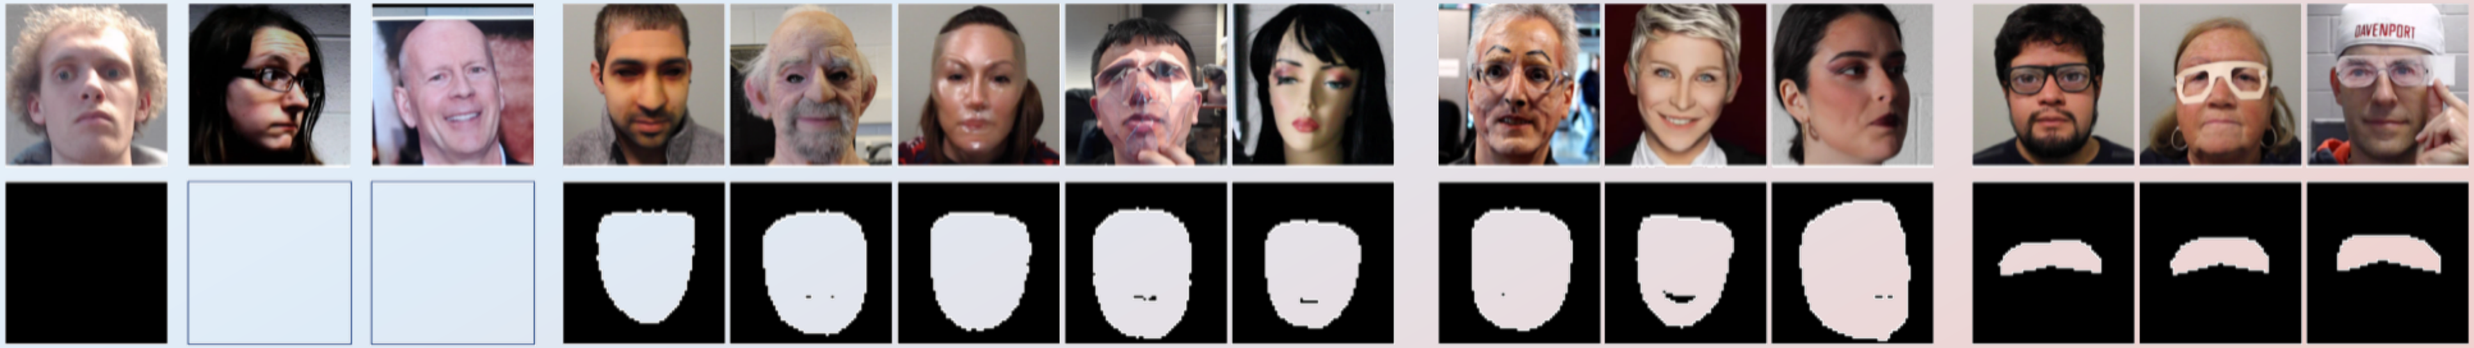

Training TRU

BTAS 2019

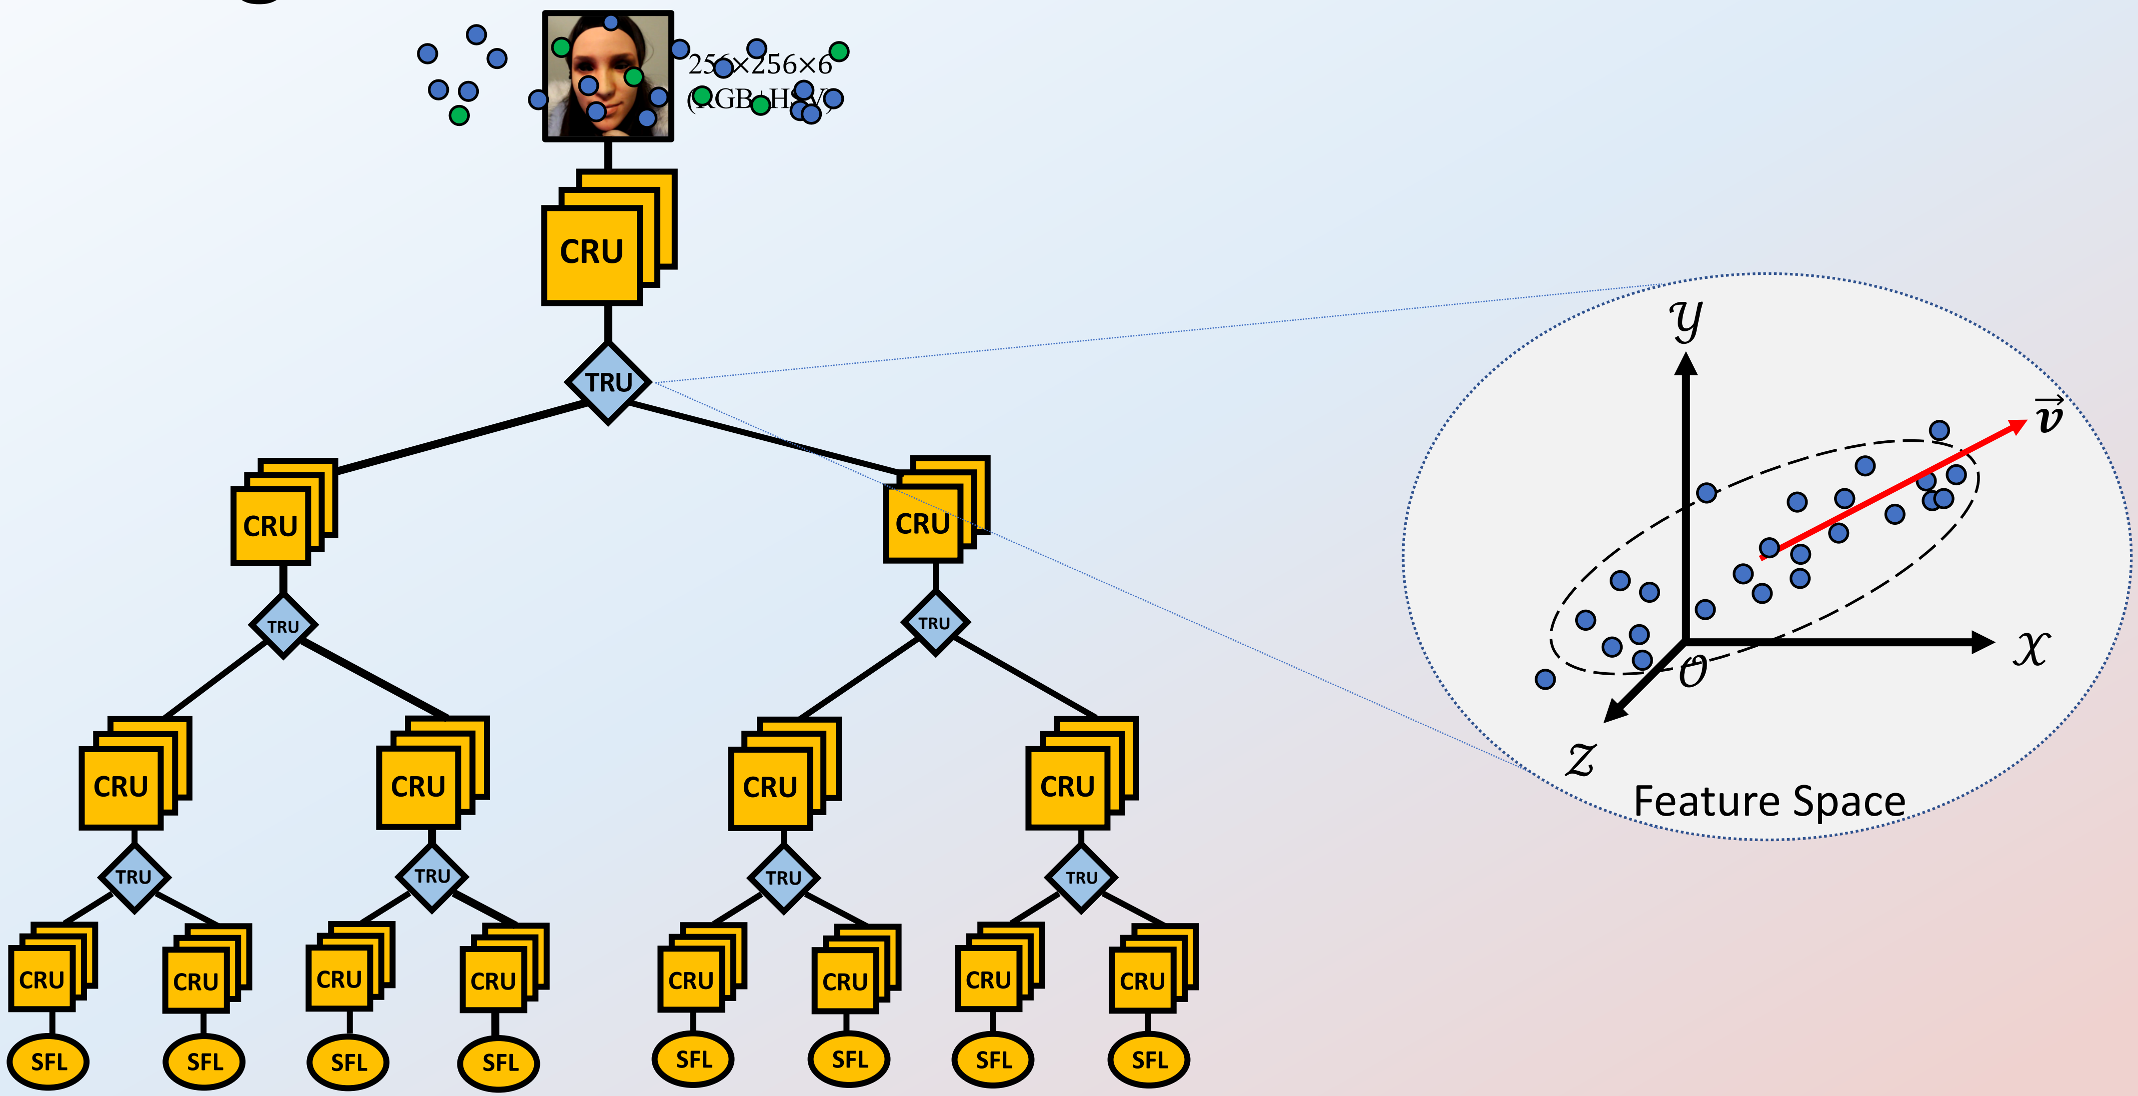

Training TRU

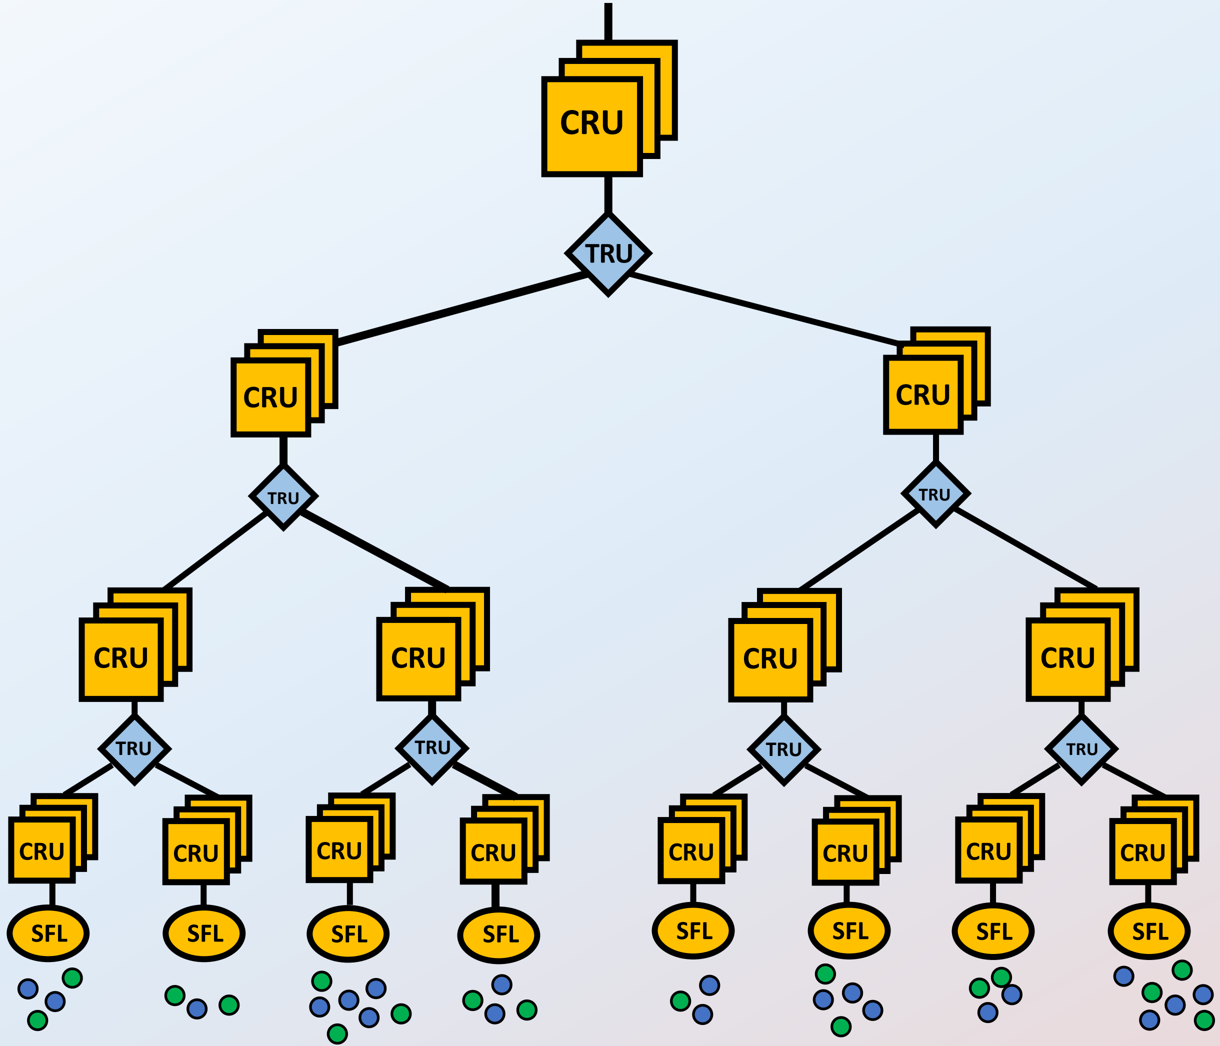

Tree Routing Unit (TRU)

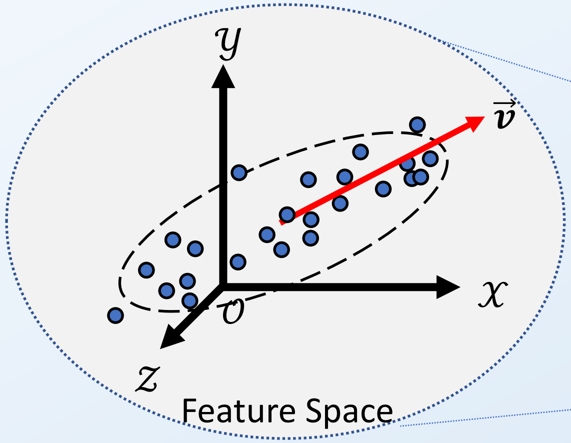

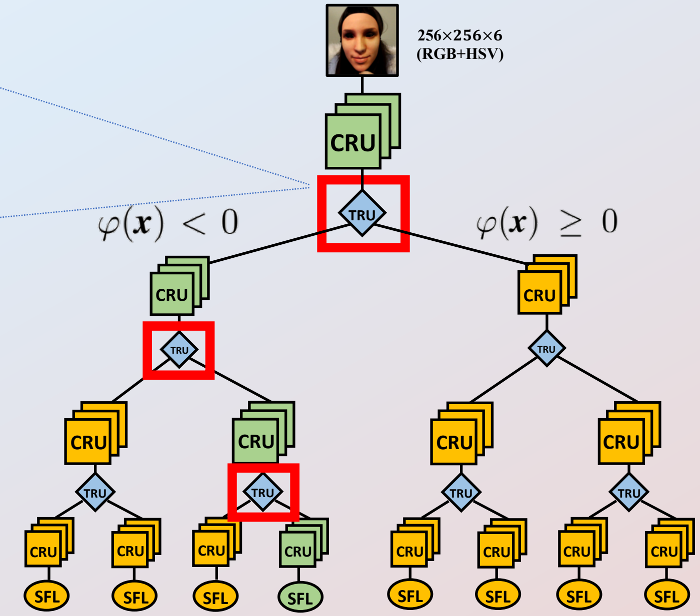

- Routing Function

$$\varphi(\boldsymbol{x}) = (\boldsymbol{x} - \boldsymbol{\mu})^T \cdot \boldsymbol{v}, \quad \|\boldsymbol{v}\| = 1$$

- Based on eigen-analysis of visiting set $\bar{\boldsymbol{X}}_{\mathcal{S}} = \boldsymbol{X}_{\mathcal{S}} - \boldsymbol{\mu}$

$$\bar{\boldsymbol{X}}_{\mathcal{S}}^T \bar{\boldsymbol{X}}_{\mathcal{S}} \boldsymbol{v} = \lambda \boldsymbol{v}$$

- We optimize:

$$\underset{\boldsymbol{v},\theta}{\arg\max}\ \lambda = \underset{\boldsymbol{v},\theta}{\arg\max}\ \boldsymbol{v}^T \bar{\boldsymbol{X}}_{\mathcal{S}}^T \bar{\boldsymbol{X}}_{\mathcal{S}} \boldsymbol{v}$$

Results

- Evaluation Metrics: ACER (the lower the better)

Methods	Replay	Print	Mask Attacks					Makeup Attacks			Partial Attacks			Avg.
			Half	Silicone	Trans.	Paper	Manne.	Obfusc.	Imperson.	Cosmetic	Funny eye	Paper Glasses	Partial Paper	
SVM+LBP[1]	20.6	18.4	31.3	21.4	45.5	11.6	13.8	59.3	23.9	16.7	35.9	39.2	11.7	26.9±14.5
Auxiliary[2]	16.8	6.9	19.3	**14.9**	52.1	8.0	12.8	55.8	13.7	**11.7**	49.0	40.5	**5.3**	23.6±18.5
Ours	**9.8**	**6.0**	**15.0**	18.7	**36.0**	**4.5**	**7.7**	**48.1**	**11.4**	14.2	**19.3**	**19.8**	8.5	**16.8±11.1**

ACER = (Spoof Error Rate (APCER) + Live Error Rate (BPCER))/2

[1] Z. Boulkenafet *et. al*. OULU-NPU: A mobile face presentation attack database with real-world variations. In FG, 2017.

[2] Y. Liu *et. al*. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In CVPR, 2018.

Results

BTAS 2019

- Evaluation Metrics: EER (the lower the better)

Methods	Replay	Print	Mask Attacks					Makeup Attacks			Partial Attacks			Avg.
			Half	Silicone	Trans.	Paper	Manne.	Obfusc.	Imperson.	Cosmetic	Funny eye	Paper Glasses	Partial Paper	
SVM+LBP	20.8	18.6	36.3	21.4	37.2	7.5	14.1	51.2	19.8	16.1	34.4	33.0	7.9	24.5±12.9
Auxiliary	14.0	4.3	**11.6**	**12.9**	**24.6**	7.8	10.0	72.3	10.1	**9.4**	21.4	**18.6**	**4.0**	17.0±17.7
Ours	**10.0**	**2.1**	14.4	18.6	26.5	**5.7**	**9.6**	**50.1**	**10.1**	13.2	**19.8**	20.5	8.8	**16.1±12.2**

[1] Z. Boulkenafet *et. al*. OULU-NPU: A mobile face presentation attack database with real-world variations. In FG, 2017.

[2] Y. Liu *et. al*. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In CVPR, 2018.

t-SNE Results

BTAS 2019

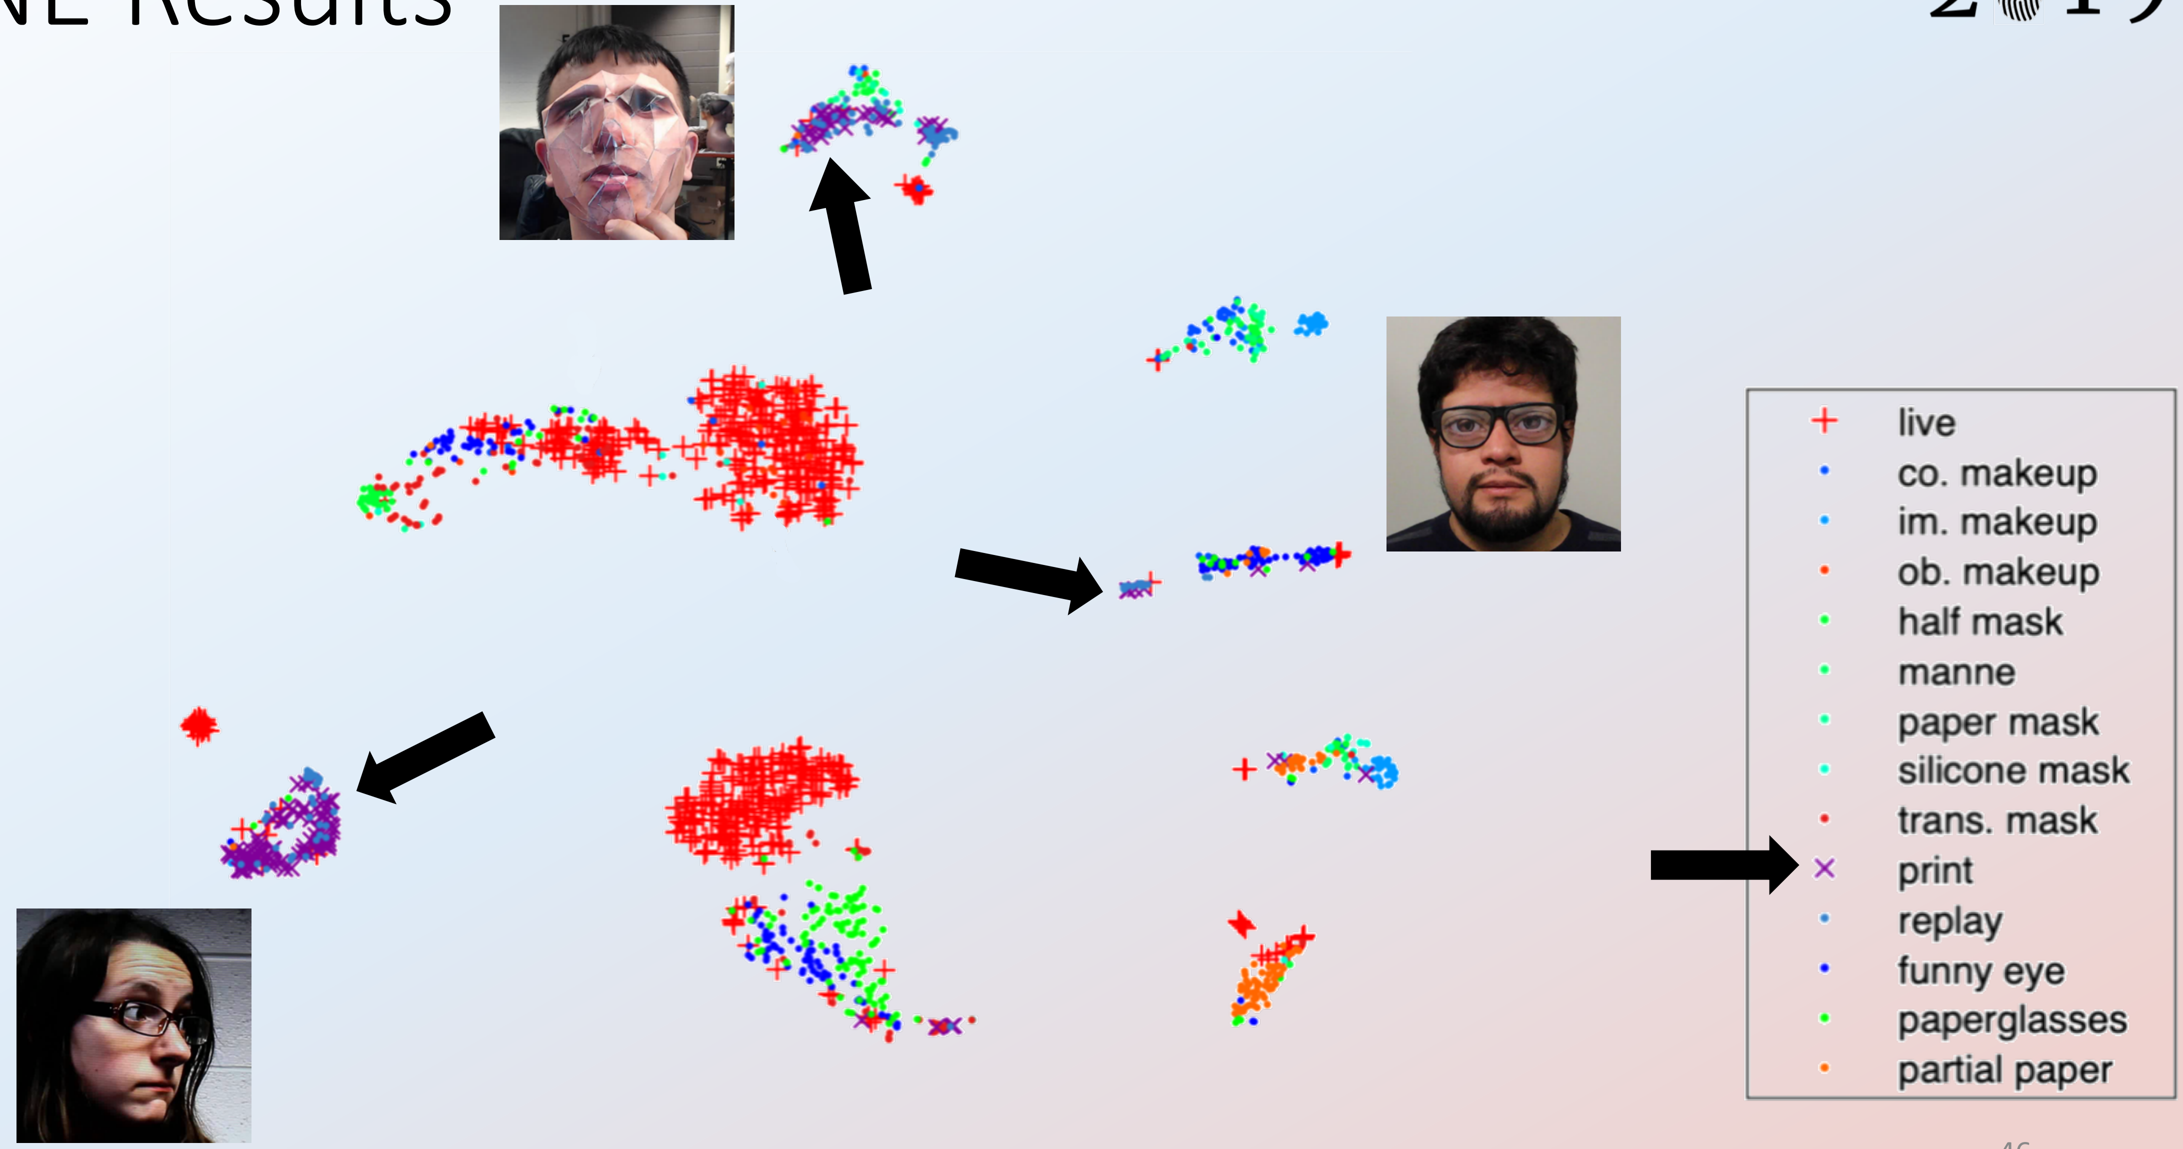

Outline

- Training-Testing difference
- Alternative/Additional Sensors
- Practical Tips
- Future

Light Reflection

- Skin and spoof material have different reflection properties:
  - Reflectance
  - 3D shape

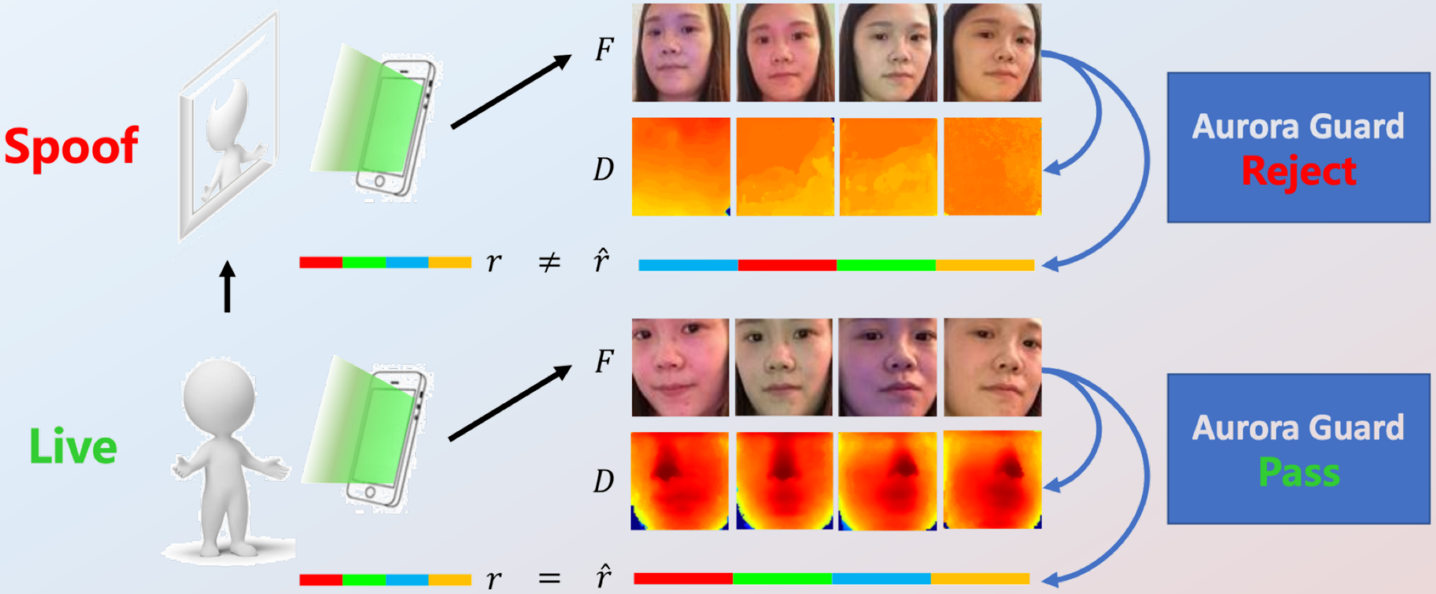

Liu et. al., Aurora guard: real-time face anti-spoofing via light reflection, arXiv 2019

Additional Sensors

- NIR
  - Human skin has different reflectance compared with spoof material
- Depth
- Thermal
- Multi-modality

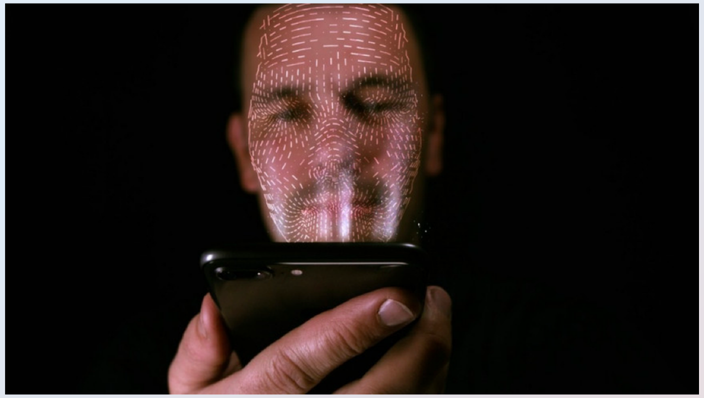

Live 3D Mask

Others

- Light field
- Polarized camera
- Structured Light
  - NIR with specific pattern (iPhone X)
- ToF (Time of flight)
  - Multi-point distance measurement

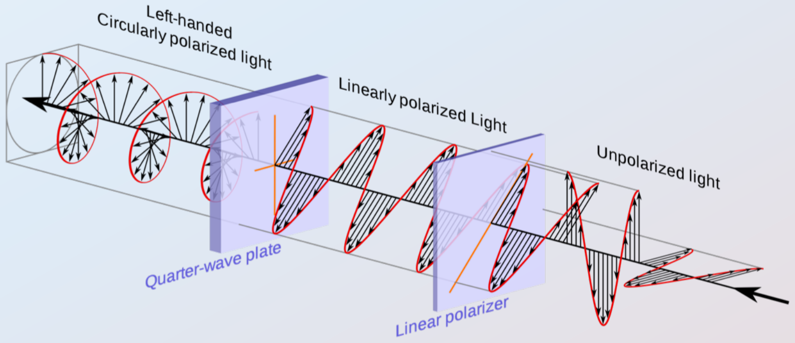

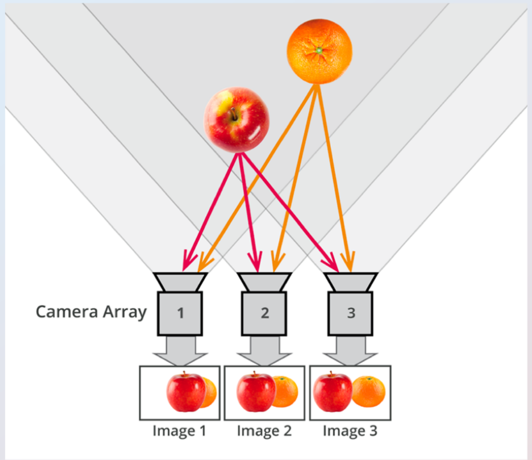

Question for Additional Sensors

- Data << RGB Data

Outline

- Training/Testing difference
- Alternative/Additional Sensors
- Practical Tips
- Future

Data are Your Friend

- More data → better performance
- Data augmentation (session II)
- (Efficient, effective) data collection

Updating Systems

- Use current model to collect failure cases
- Add failure cases to training set to fine-tune the model
- Update the current model
- Repeat several times

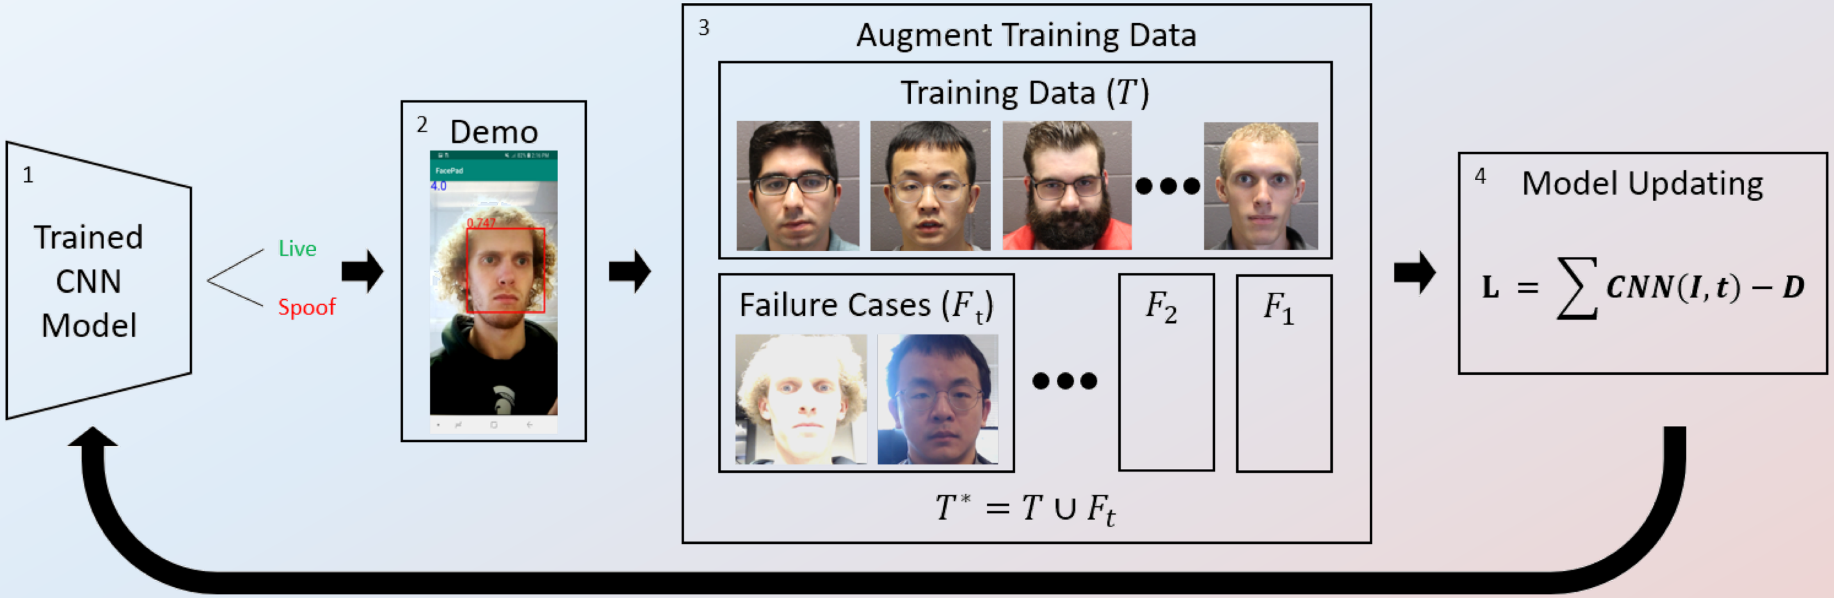

Updating Systems

- Manage the training data, not just mix everything
  - Eg. Base data 80%, New data 20%
  - Add subclasses based on lighting, walking and etc

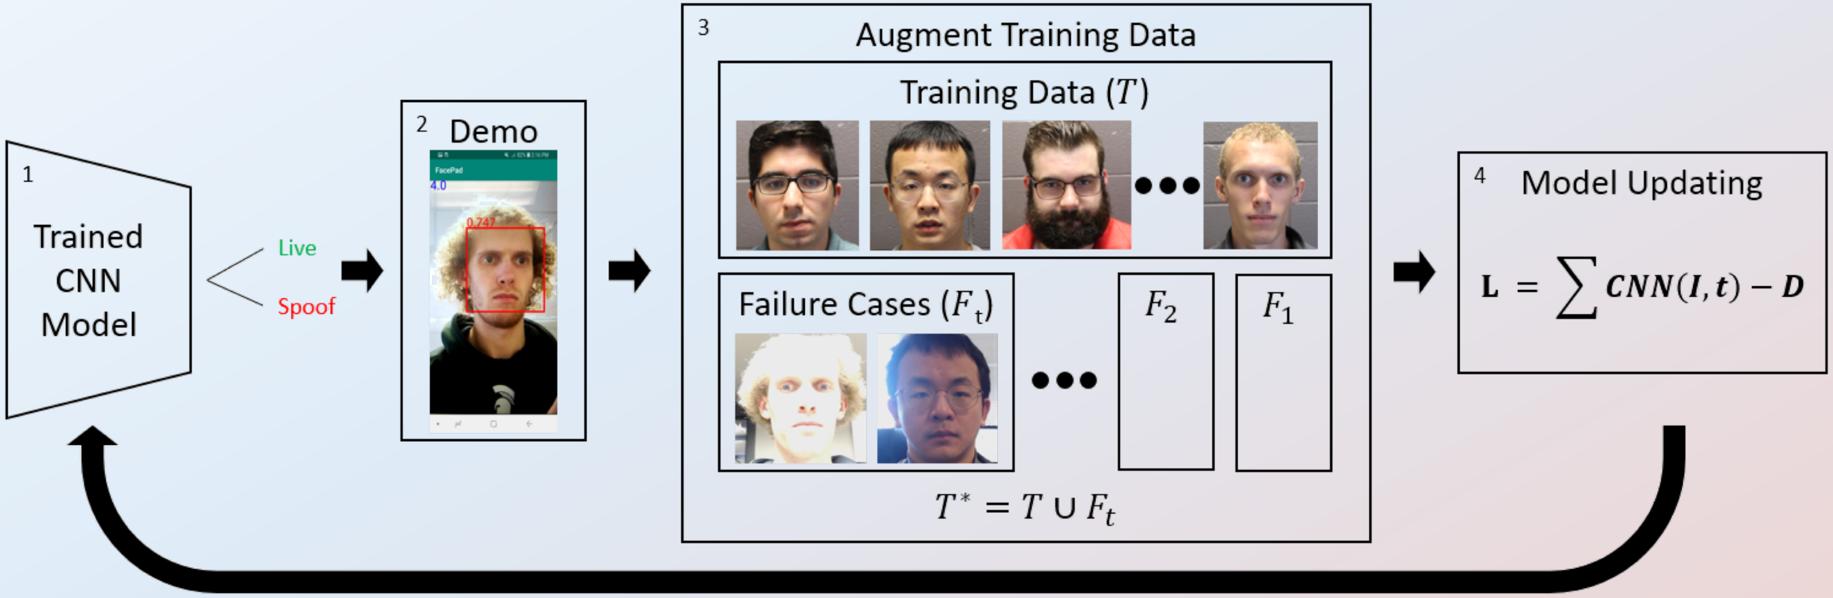

Image Quality is the Devil

BTAS 2019

- Image resolution
- JPEG compression
  - Check the image bitrate
- Dark environment → ISO noise

Image Quality is the Devil

- Image resolution
- JPEG compression
    - Check the image bitrate
- Dark environment → ISO noise

Outline

- Training/Testing difference
- Alternative/Additional Sensors
- Practical Tips
- Summary and Future

Unsolved Problems

BTAS 2019

- Training/Testing difference
- Explainablity
- New attacks
- Unknown attack
- Data and evaluation

Problem 1: Training-Testing Difference

BTAS 2019

- Cross-database testing performances are still poor
  - EER for intra-testing: ~ 0% – 5%
  - EER for inter-testing: ~ 15% - 50%
- Can we use few-shot learning to improve the cross-database testing?

Problem 2: Explainablity

- Spatial explainablity
- Temporal explainablity
- Spoofing process explainablity
- Research on camera and imaging

Problem 3: New Attacks

- Makeup attacks
- Counter attacks to current methods
  - 3D mask attacks with flashing light → rPPG methods
  - Adversarial attacks → Texture based methods

Problem 4: Unknown Attacks

BTAS 2019

- Similar situation to cross-database testing
- Can we leverage the knowledge from other unknown object detection tasks?
- Identity variations > anti-spoofing variation

Problem 5: Data and Evaluation

- Intra-testing protocols too easy
- Inter-testing protocols too hard
- Represent previous problems as the testing protocols

Summary

- What and why face anti-spoofing?
- Traditional methods
- Deep learning methods
- Unknown attacks
- Additional sensors
- Practical tips