BIOTAMPERNET: AFFINITY-GUIDED STATE-SPACE MODEL DETECTING TAMPERED BIOMEDICAL IMAGES

SOUMYAROOP NANDI[1,2], **PREM NATARAJAN**[1,2,3]
[1]USC INFORMATION SCIENCES INSTITUTE, MARINA DEL REY, CA, USA
[2]USC THOMAS LORD CS DEPARTMENT, LOS ANGELES, CA, USA
[3]CAPITAL ONE
{SOUMYARN, PREMKUMN}@USC.EDU

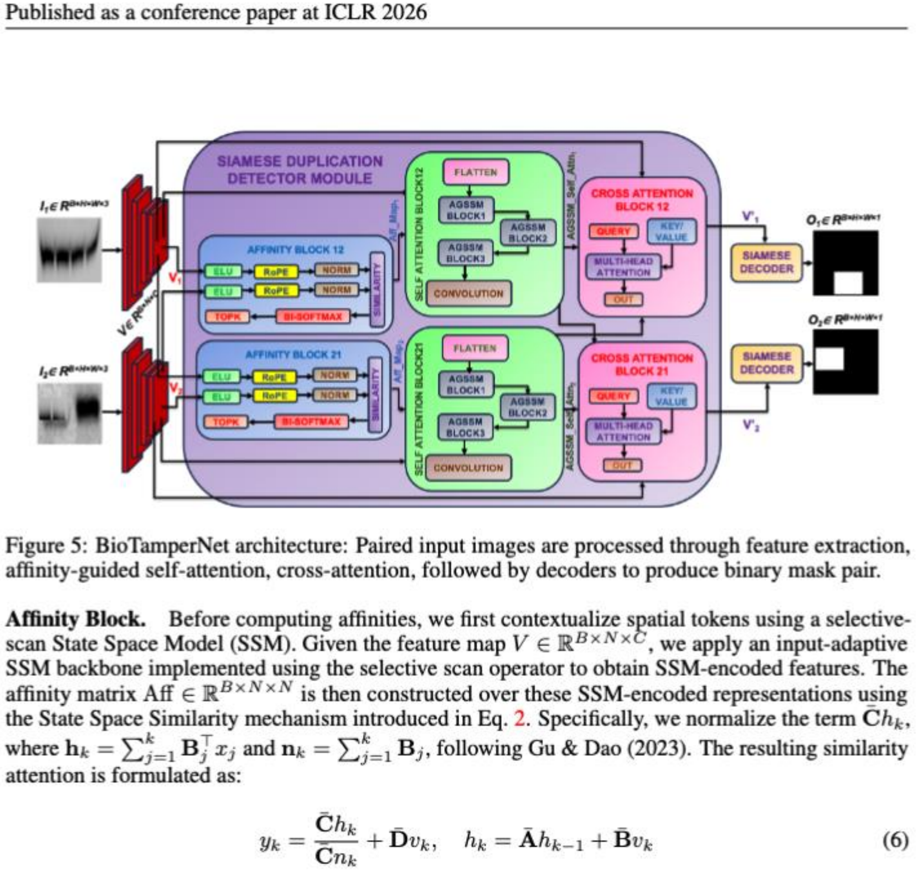

Published as a conference paper at ICLR 2026

Figure 5: BioTamperNet architecture: Paired input images are processed through feature extraction, affinity-guided self-attention, cross-attention, followed by decoders to produce binary mask pair.

Affinity Block. Before computing affinities, we first contextualize spatial tokens using a selective-scan State Space Model (SSM). Given the feature map $V \in \mathbb{R}^{B\times N\times C}$, we apply an input-adaptive SSM backbone implemented using the selective scan operator to obtain SSM-encoded features. The affinity matrix $\text{Aff} \in \mathbb{R}^{B\times N\times N}$ is then constructed over these SSM-encoded representations using the State Space Similarity mechanism introduced in Eq. 2. Specifically, we normalize the term $\bar{\mathbf{C}}h_k$, where $\mathbf{h}_k = \sum_{j=1}^{k} \mathbf{B}_j^\top x_j$ and $\mathbf{n}_k = \sum_{j=1}^{k} \mathbf{B}_j$, following Gu & Dao (2023). The resulting similarity attention is formulated as:

$$y_k = \frac{\bar{\mathbf{C}}h_k}{\bar{\mathbf{C}}n_k} + \bar{\mathbf{D}}v_k, \quad h_k = \bar{\mathbf{A}}h_{k-1} + \bar{\mathbf{B}}v_k \tag{6}$$

ROADMAP

- **What is Scientific Image Misconduct?**
    - Bioforensic Image Types
    - Bioforensic Duplication Detection Tasks
    - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
    - BioTamperNet Contributions
    - BioTamperNet Training Setup
    - State Space Model Overview
    - BioTamperNet Architecture
    - Experimental Evaluation and Results
    - Directions to Improve

IMPACT OF SCIENTIFIC MISCONDUCT

- **Unethical**: Negatively impact reputation and erodes society's trust in science
- **Reproducibility**: Results from fraudulent papers are often not reproducible [*Miyakawa et al. Molecular Brain, 2020*]
- **Retractions**: Out of 200 papers surveyed, 41 corrections and 5 retractions [*Bik et al. Molecular and Cellular Biology, 2018*]
- **Financial Loss**: Financial cost is associated with retractions [*Stern et al. Elife, 3:e02956, 2014*]
- **Generative AI**: Highly accurate representations of complex scientific visualizations [*Kim et al. Generative AI can fabricate advanced scientific visualizations: ethical implications and strategic mitigation framework, AI and Ethics 2024*]

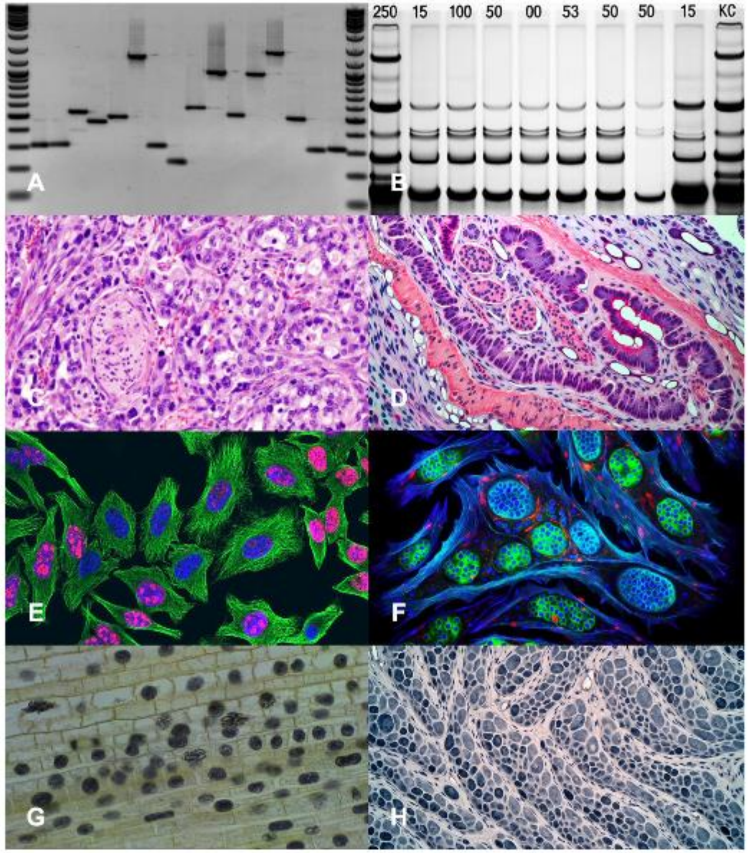

Fig 16: Kim et al. AI and Ethics 2024 - Laboratory images compared to images generated by DALL-E 3

WHY REVIEWERS CANNOT DETECT?

- **Scale**: Image duplication has increased from ~2% to 4% since 2003 [*Bik et al. Molecular Biology, 2016*]
- **Complexity**: Complex manipulations may span across documents and beyond human ability [*Bik et al. Molecular Biology, 2016*]
- **Abstract Patterns**: Difficult for human to analyze random patterns
- **Diverse Manipulation**: Difficult for reviewers to be aware of a wide range of existing manipulations

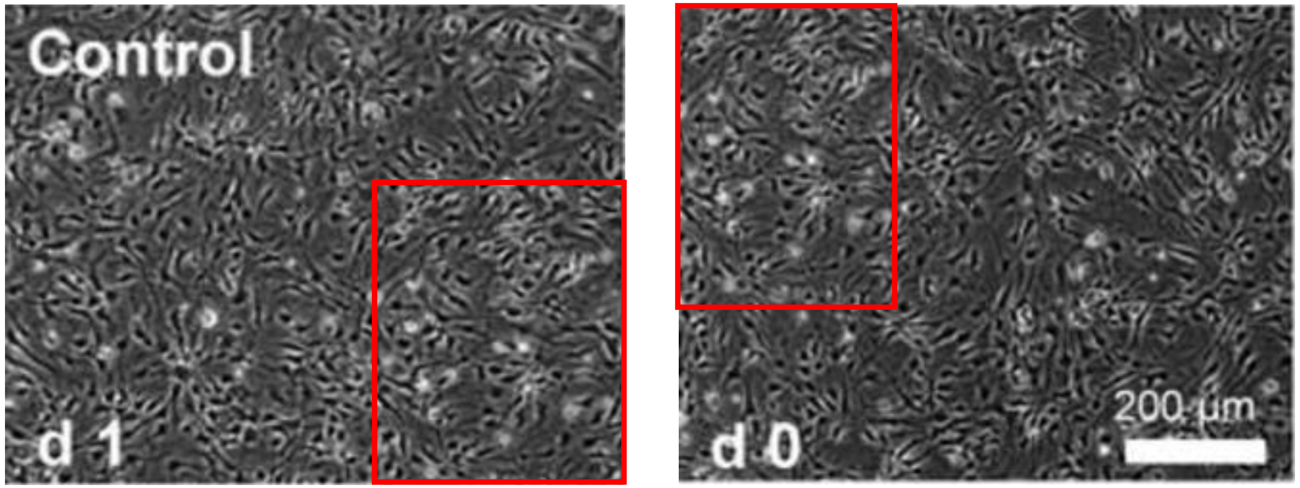

Fig 17: Esfandiari et. al., SCD 2012

ROADMAP

- What is Scientific Image Misconduct?
  - **Bioforensic Image Types**
  - Bioforensic Duplication Detection Tasks
  - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - BioTamperNet Contributions
  - BioTamperNet Training Setup
  - State Space Model Overview
  - BioTamperNet Architecture
  - Experimental Evaluation and Results
  - Directions to Improve

BIOFORENSIC IMAGE TYPES - MICROSCOPY

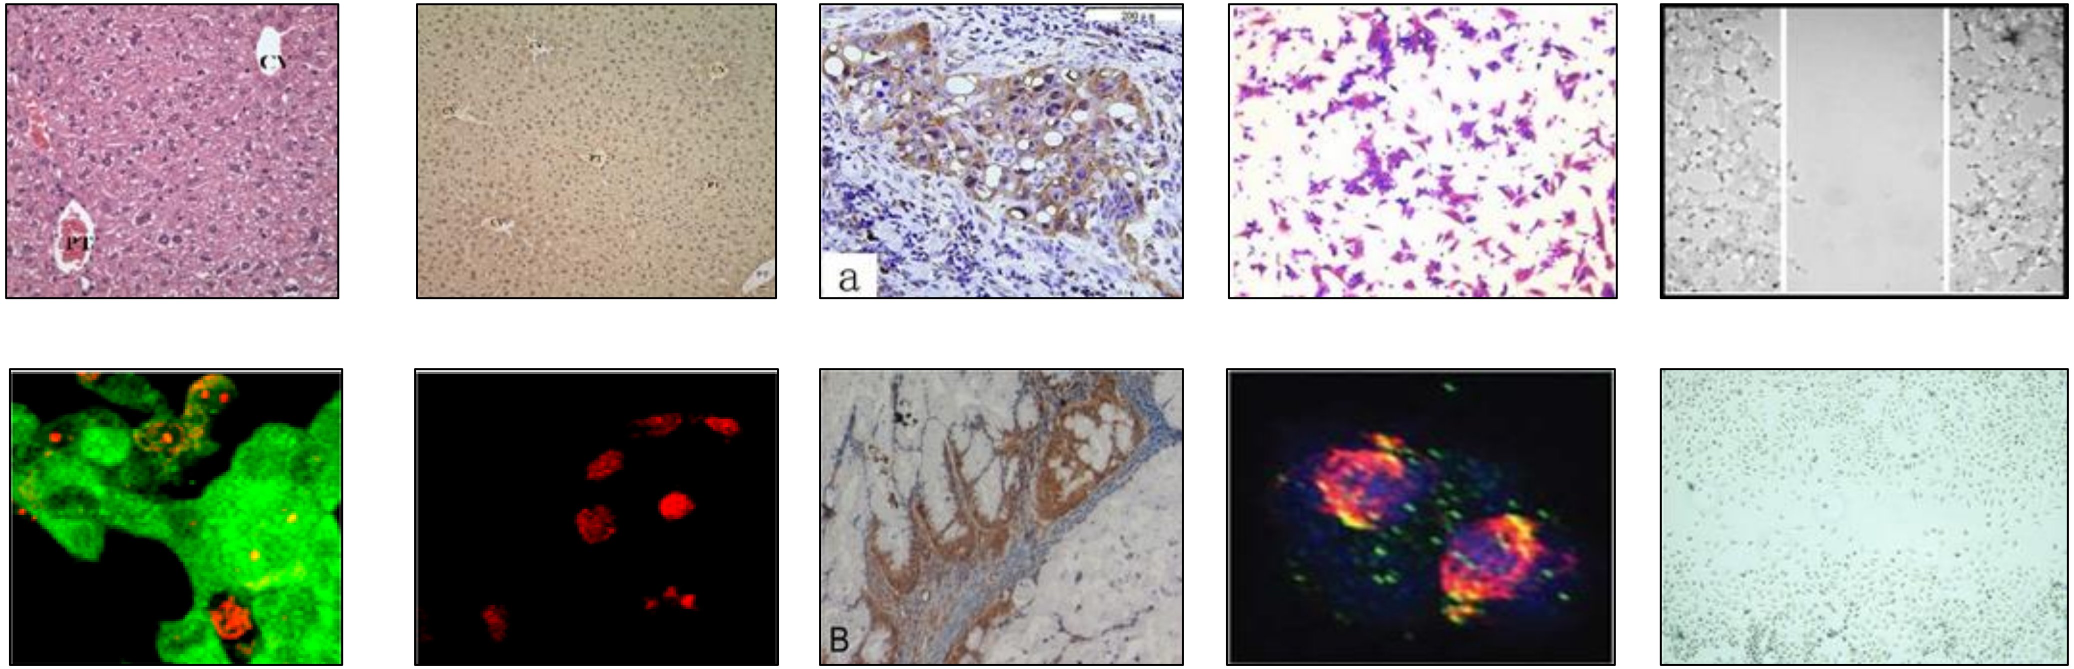

Usually Cell/Tissue Images Taken Under Microscope

BIOFORENSIC IMAGE TYPES - BLOT/GEL

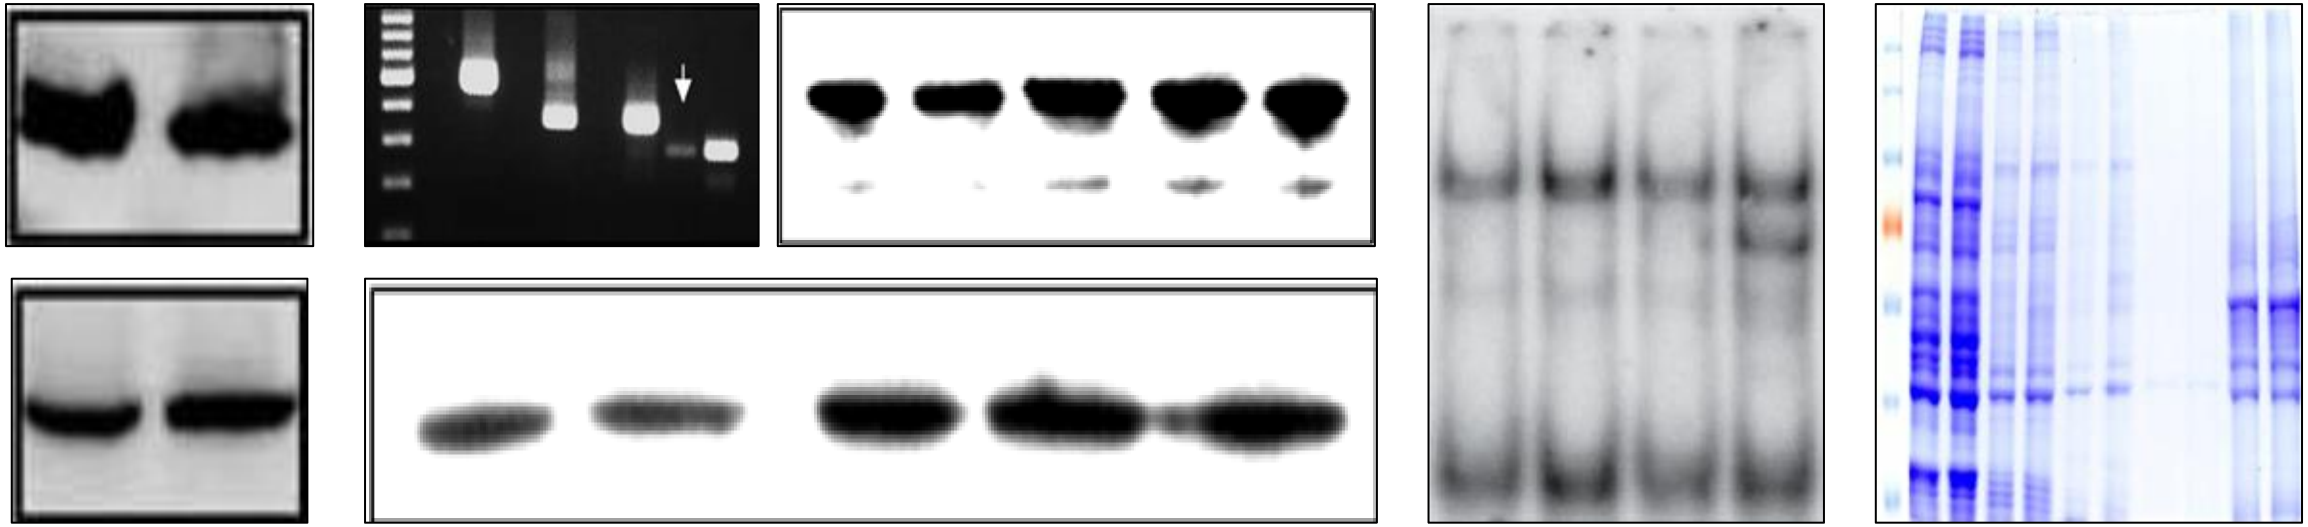

Used for Protein RNA and DNA Analysis :
Western, Northern, Southern Blots and Gels

BIOFORENSIC IMAGE TYPES - FACS

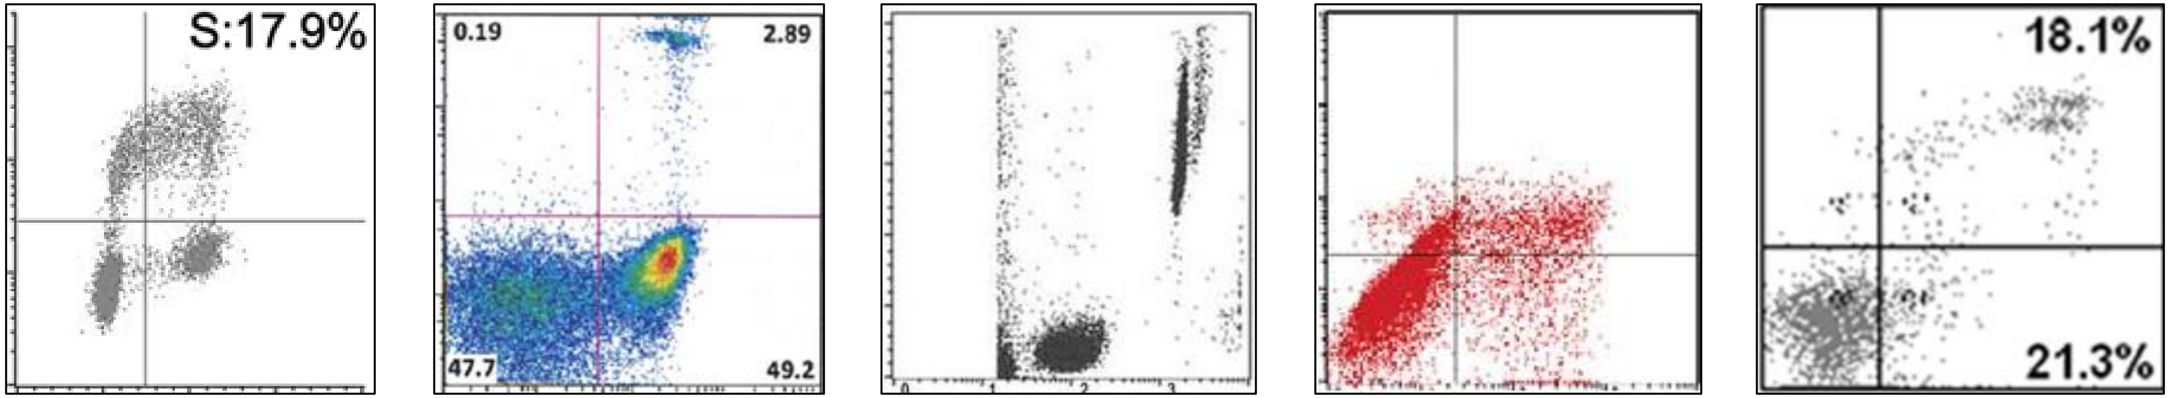

Usually Flow Cytometry Images, similar to Scatter Plots

BIOFORENSIC IMAGE TYPES - MACROSCOPY

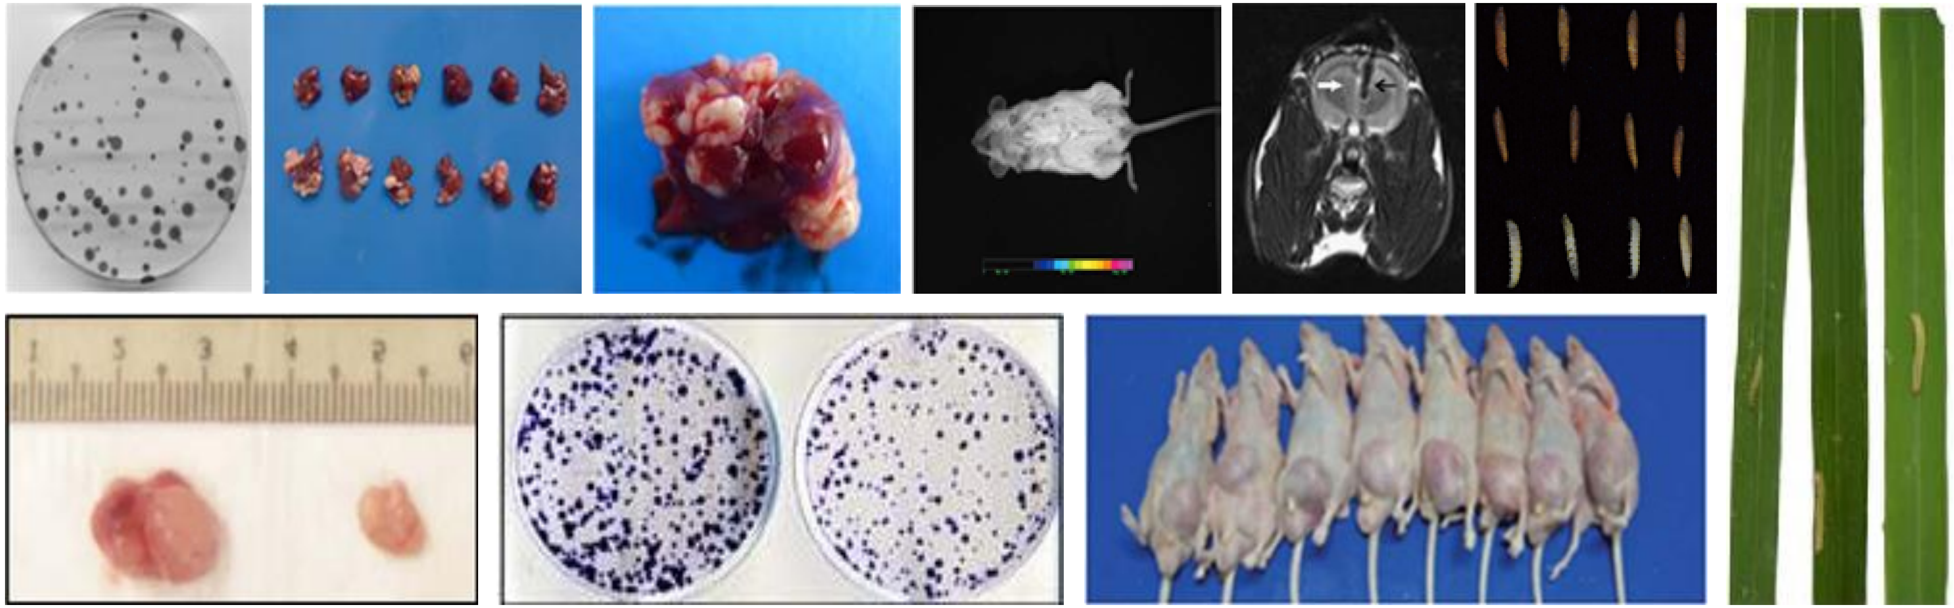

Miscellaneous Images like Scans, Leaves, Tissues etc.

ROADMAP

- What is Scientific Image Misconduct?
  - Bioforensic Image Types
  - **Bioforensic Duplication Detection Tasks**
  - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - BioTamperNet Contributions
  - BioTamperNet Training Setup
  - State Space Model Overview
  - BioTamperNet Architecture
  - Experimental Evaluation and Results
  - Directions to Improve

BIOFORENSICS FORGERY TASK –
EXTERNAL DUPLICATION DETECTION (EDD)

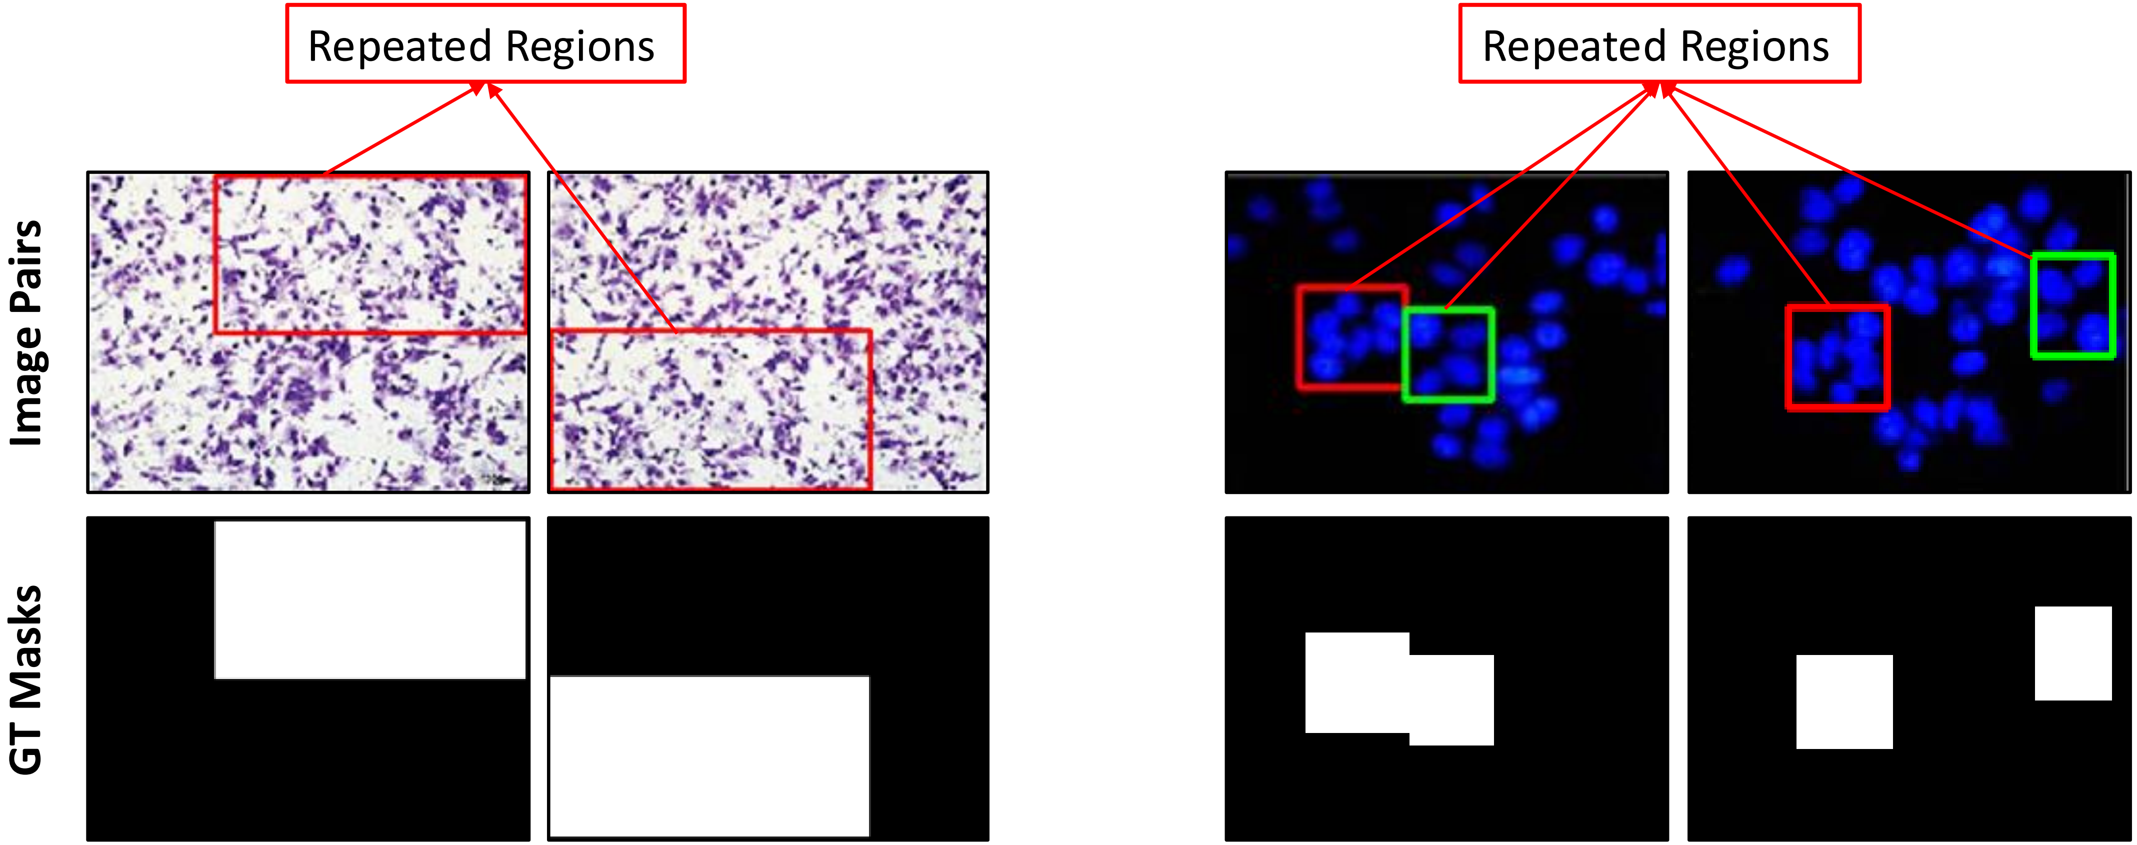

Repeated Regions

Repeated Regions

BIOFORENSICS FORGERY TASK – INTERNAL DUPLICATION DETECTION (IDD)

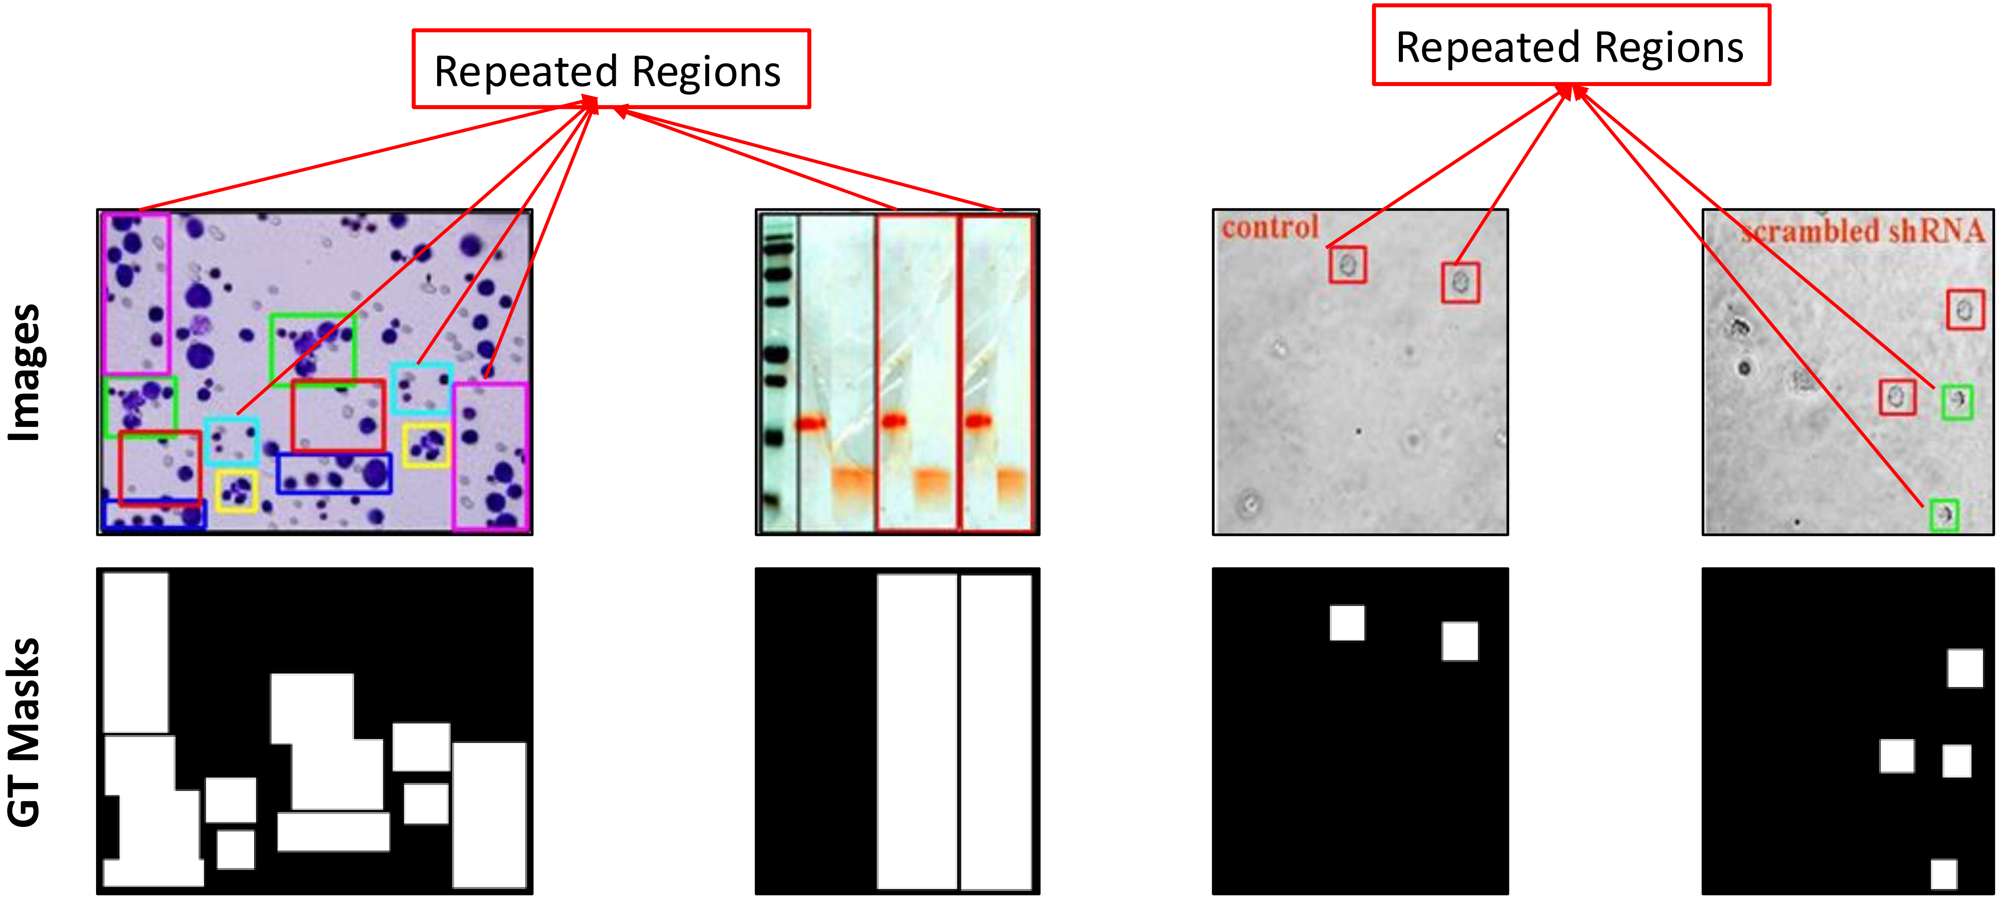

Repeated Regions

Repeated Regions

Images

GT Masks

BIOFORENSICS FORGERY TASK –
CUT/SHARP TRANSITION DETECTION (CSTD)

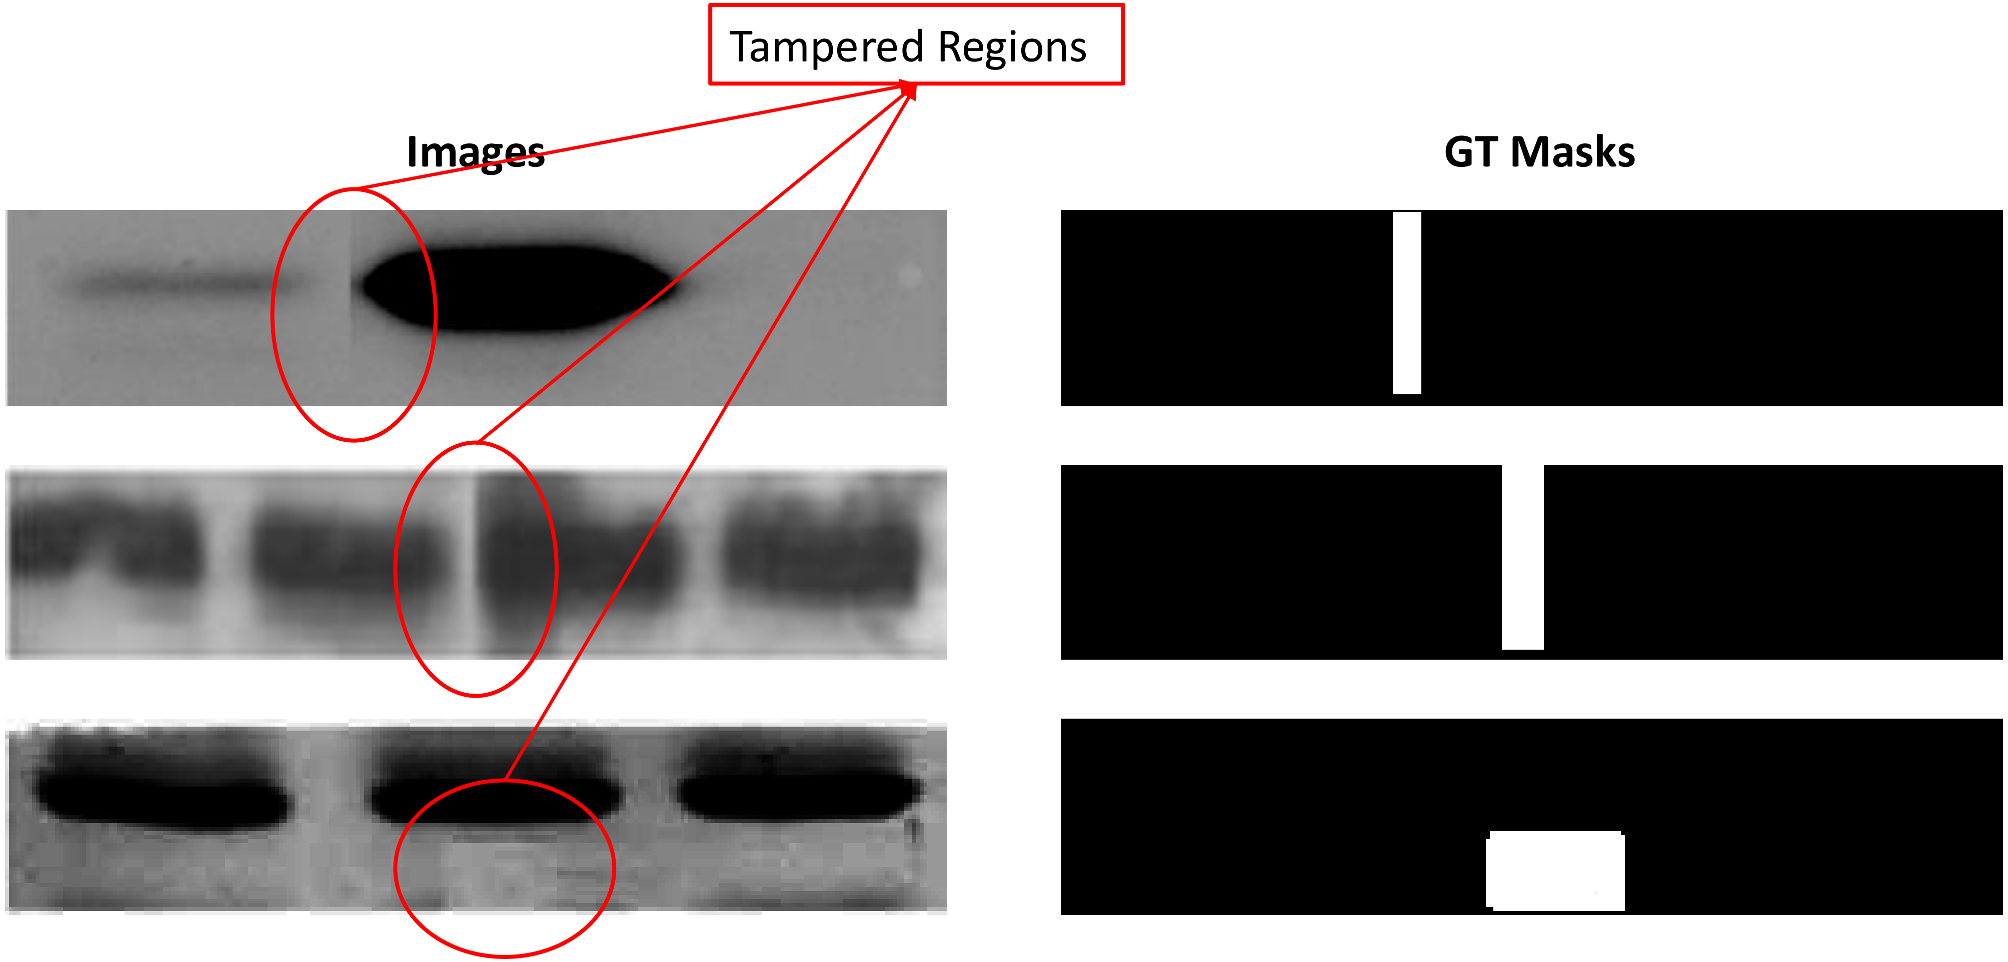

ROADMAP

- What is Scientific Image Misconduct?
  - Bioforensic Image Types
  - Bioforensic Duplication Detection Tasks
  - **Why Biomedical Forgery Detection is Hard?**
- Solution Proposed – BioTamperNet Model
  - BioTamperNet Contributions
  - BioTamperNet Training Setup
  - State Space Model Overview
  - BioTamperNet Architecture
  - Experimental Evaluation and Results
  - Directions to Improve

CHALLENGES IN BIOMEDICAL FORENSICS - IMAGE ARTIFACTS

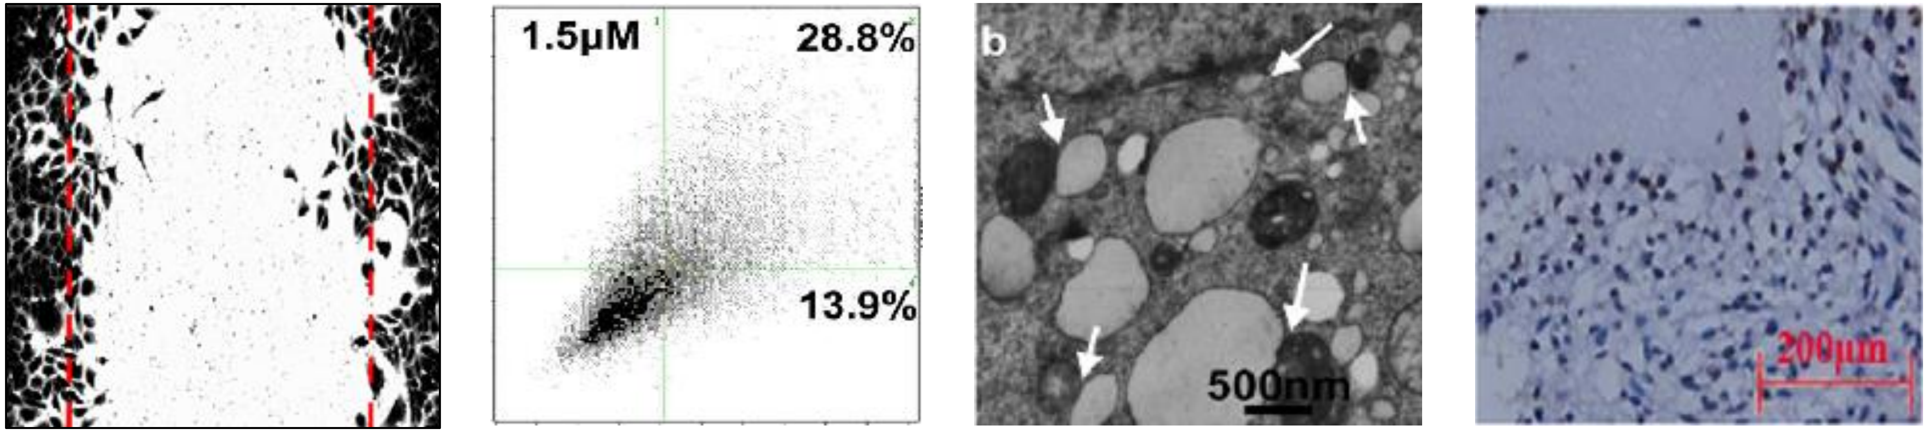

Dotted Lines (a)

Text (b)

Arrows (c)

Scale (d)

Fig 18: Examples of annotation artifacts in biomedical images
(a) Dotted lines (b) alphanumeric text (c) arrows (d) scale

CHALLENGES IN BIOMEDICAL FORENSICS - FIGURE SEMANTICS

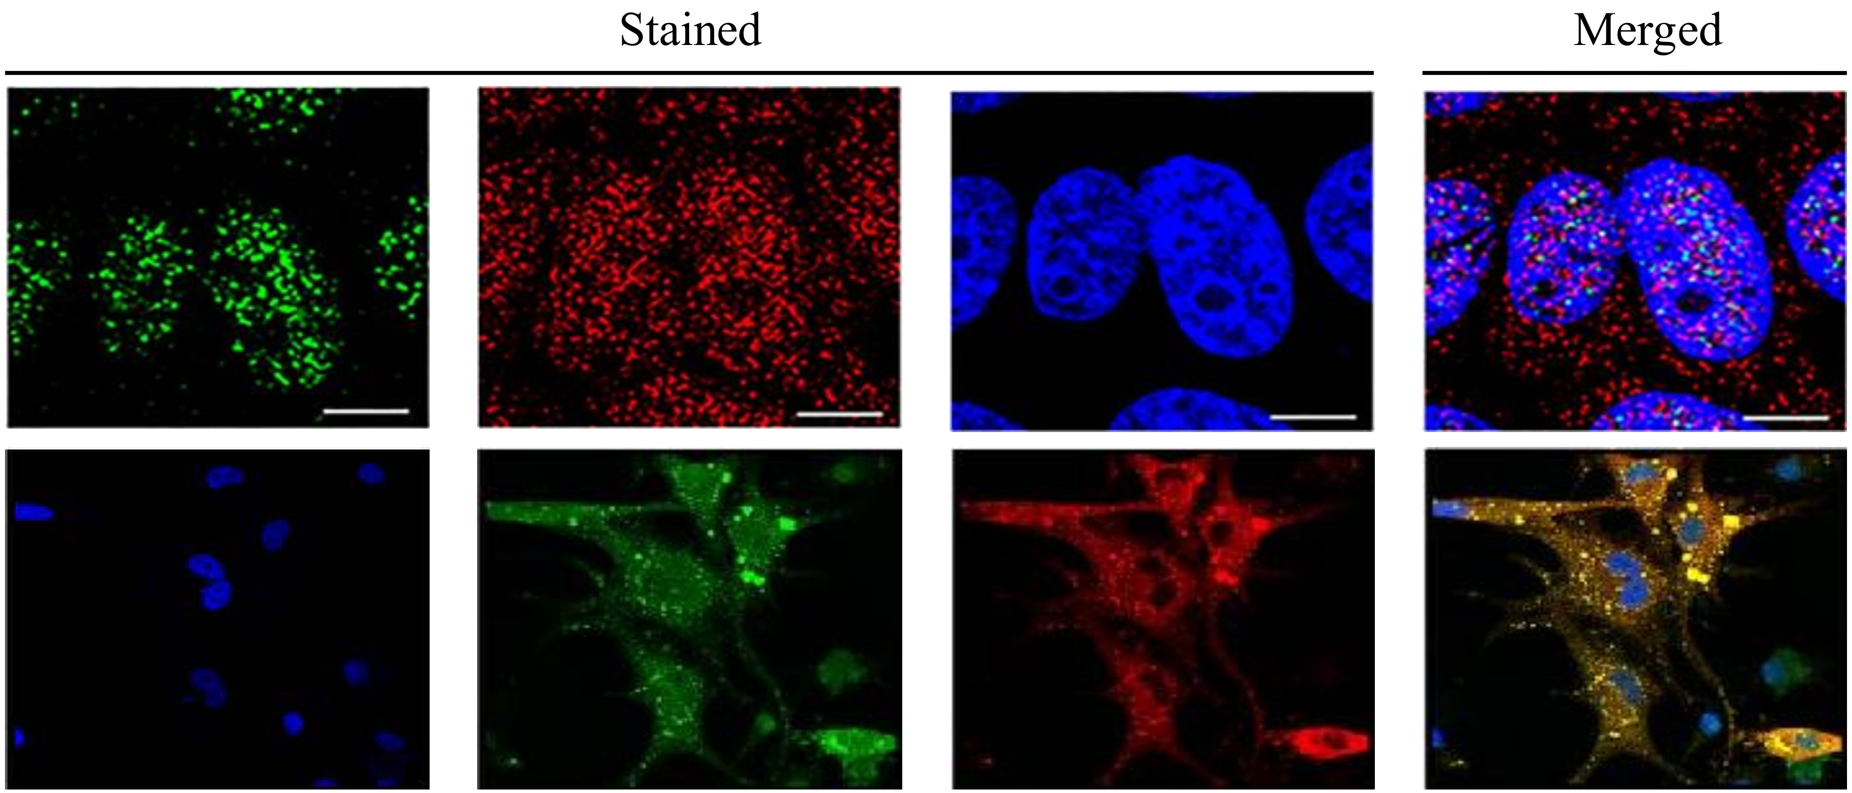

Fig 19: Staining of Microscopy Images

CHALLENGES IN BIOMEDICAL FORENSICS - FIGURE SEMANTICS

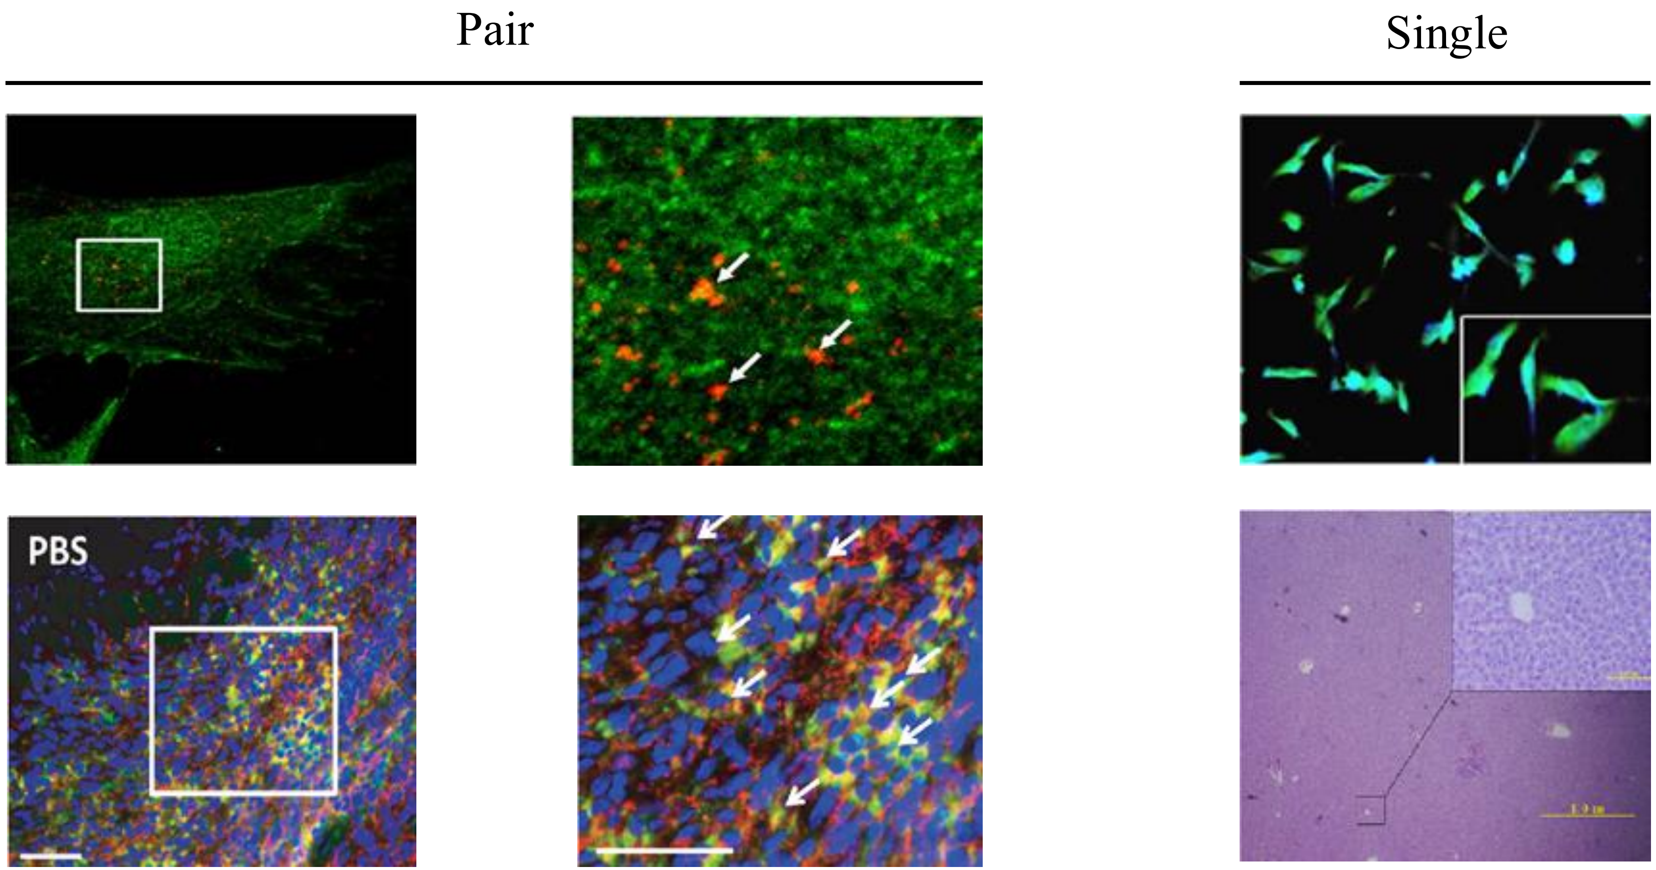

Fig 20: Pairs of Original and Zoomed Regions in Images

CHALLENGES IN BIOMEDICAL FORENSICS - IMAGE TEXTURE

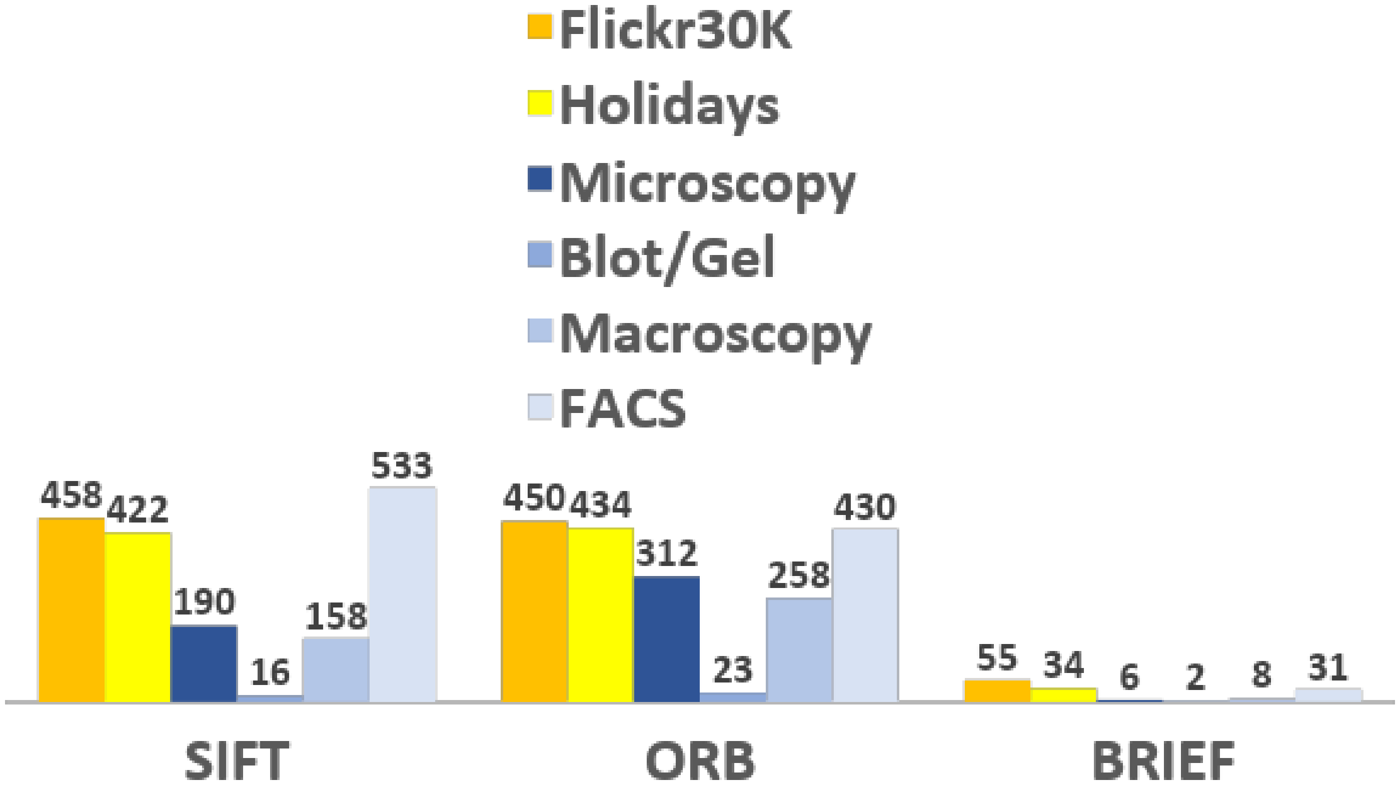

Fig 21: Plain or Pattern-like Texture in Biomedical Images showing Median number of Identified Key-points

CHALLENGES IN BIOMEDICAL FORENSICS - HARD NEGATIVES

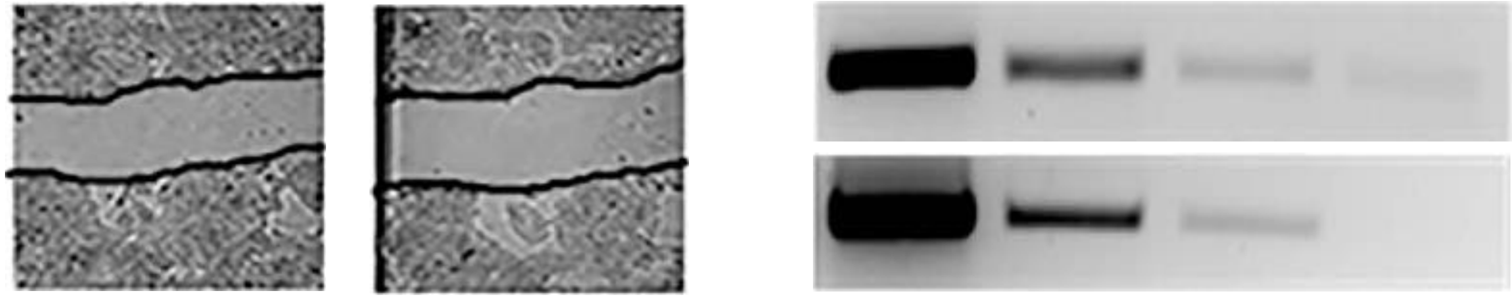
Microscopy

Blot/Gel

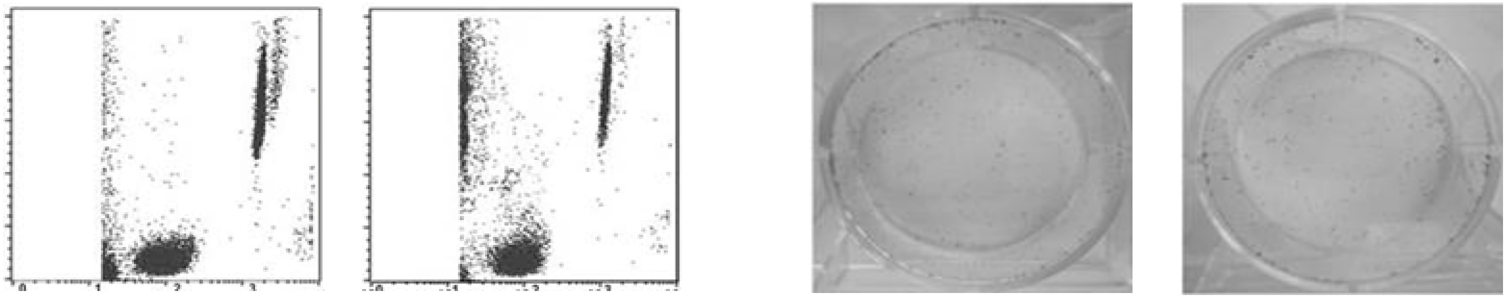
FACS

Macroscopy

Fig 22: Similar looking Hard Negative Samples From Blot/Gel, Macroscopy, FACS and Microscopy classes in clockwise order

ROADMAP

- What is Scientific Image Misconduct?
  - Bioforensic Image Types
  - Bioforensic Duplication Detection Tasks
  - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - **BioTamperNet Contributions**
  - BioTamperNet Training Setup
  - State Space Model Overview
  - BioTamperNet Architecture
  - Experimental Evaluation and Results
  - Directions to Improve

BIOTAMPERNET CONTRIBUTIONS

- **BioTamperNet -** A unified framework for paired and single-image biomedical forgery localization
- **Affinity-guided linear attention -** Lightweight SSM-inspired self and cross-attention for structured duplication detection
- **Synthetic + GAN-based augmentation** to address data scarcity in biomedical forensics
- **State-of-the-art performance** on the BioFors (Sabir et al., 2021) real-world retracted benchmark

ROADMAP

- What is Scientific Image Misconduct?
  - Bioforensic Image Types
  - Bioforensic Duplication Detection Tasks
  - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - BioTamperNet Contributions
  - **BioTamperNet Training Setup**
  - State Space Model Overview
  - BioTamperNet Architecture
  - Experimental Evaluation and Results
  - Directions to Improve

BIOTAMPERNET TRAINING SETUP - EDD

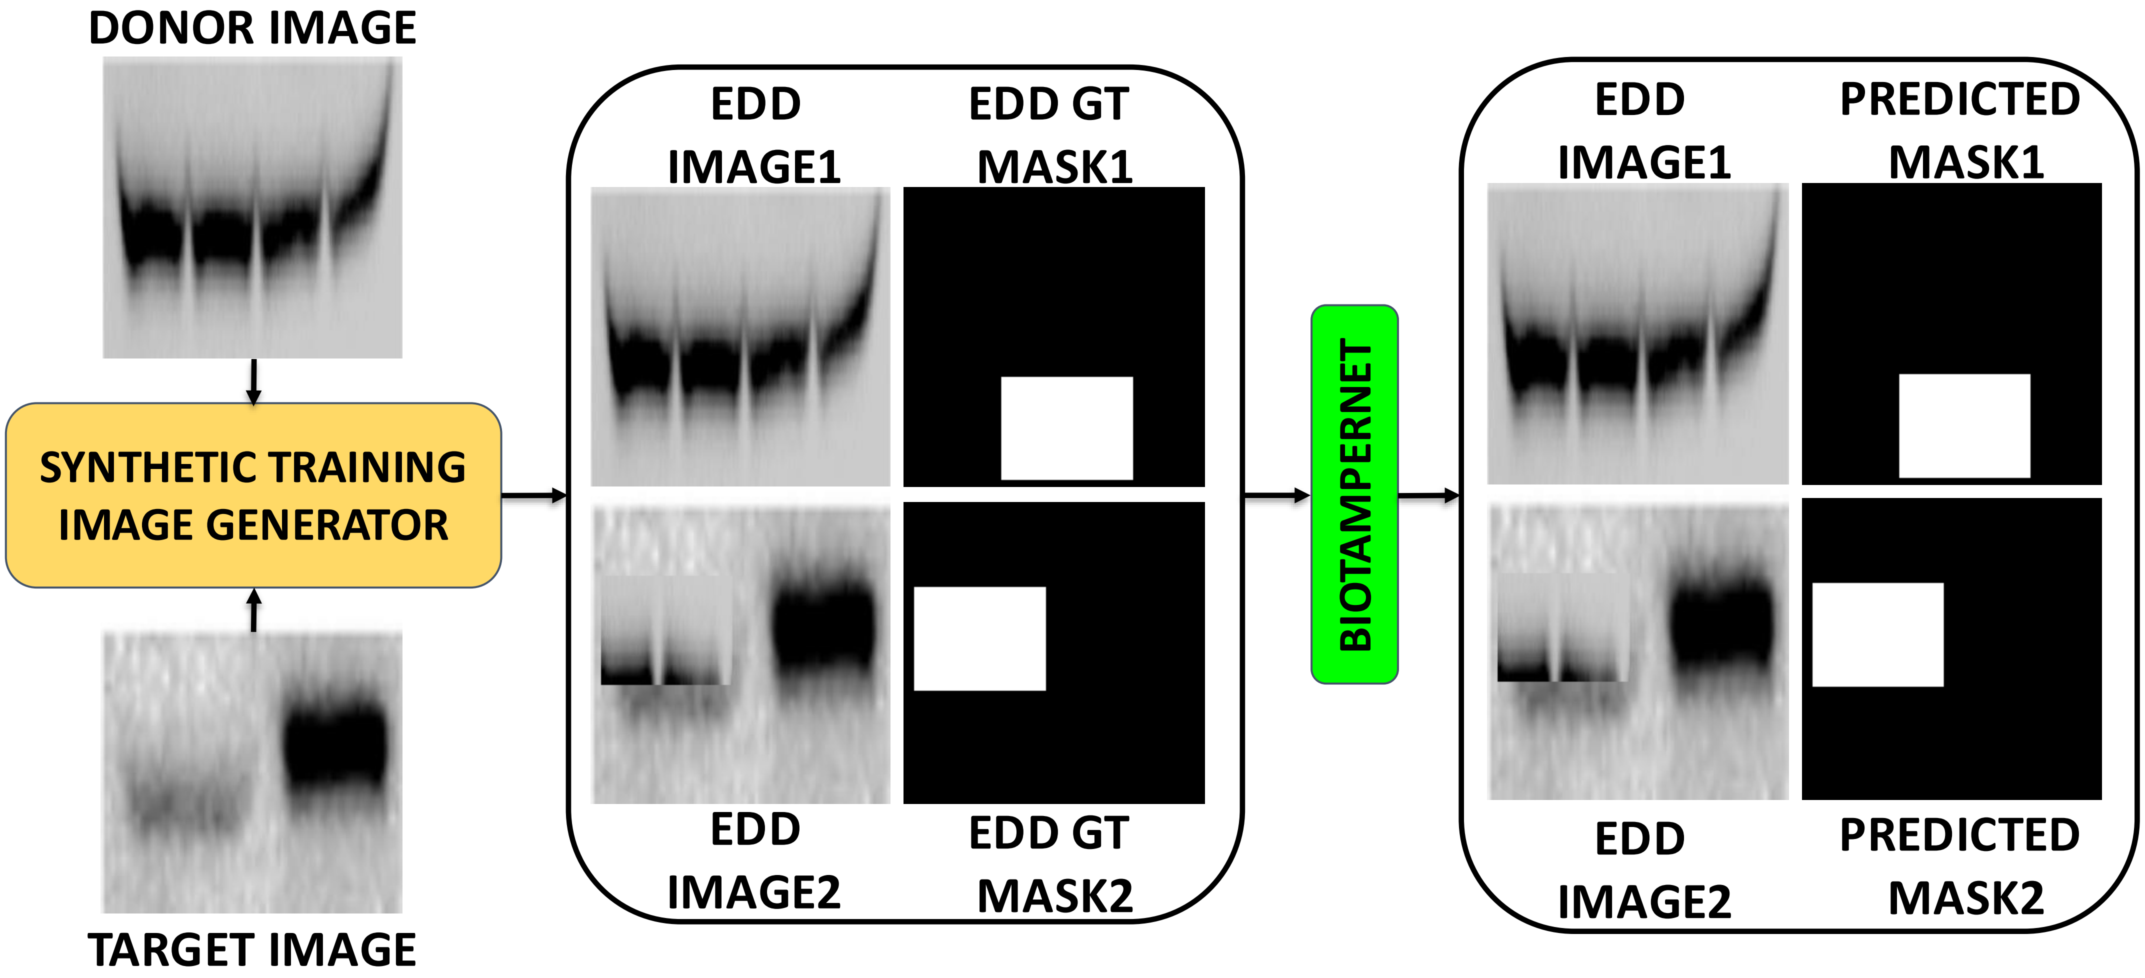

DONOR IMAGE

SYNTHETIC TRAINING IMAGE GENERATOR

TARGET IMAGE

EDD IMAGE1

EDD GT MASK1

EDD IMAGE2

EDD GT MASK2

BIOTAMPERNET

EDD IMAGE1

PREDICTED MASK1

EDD IMAGE2

PREDICTED MASK2

BIOTAMPERNET TRAINING SETUP - IDD

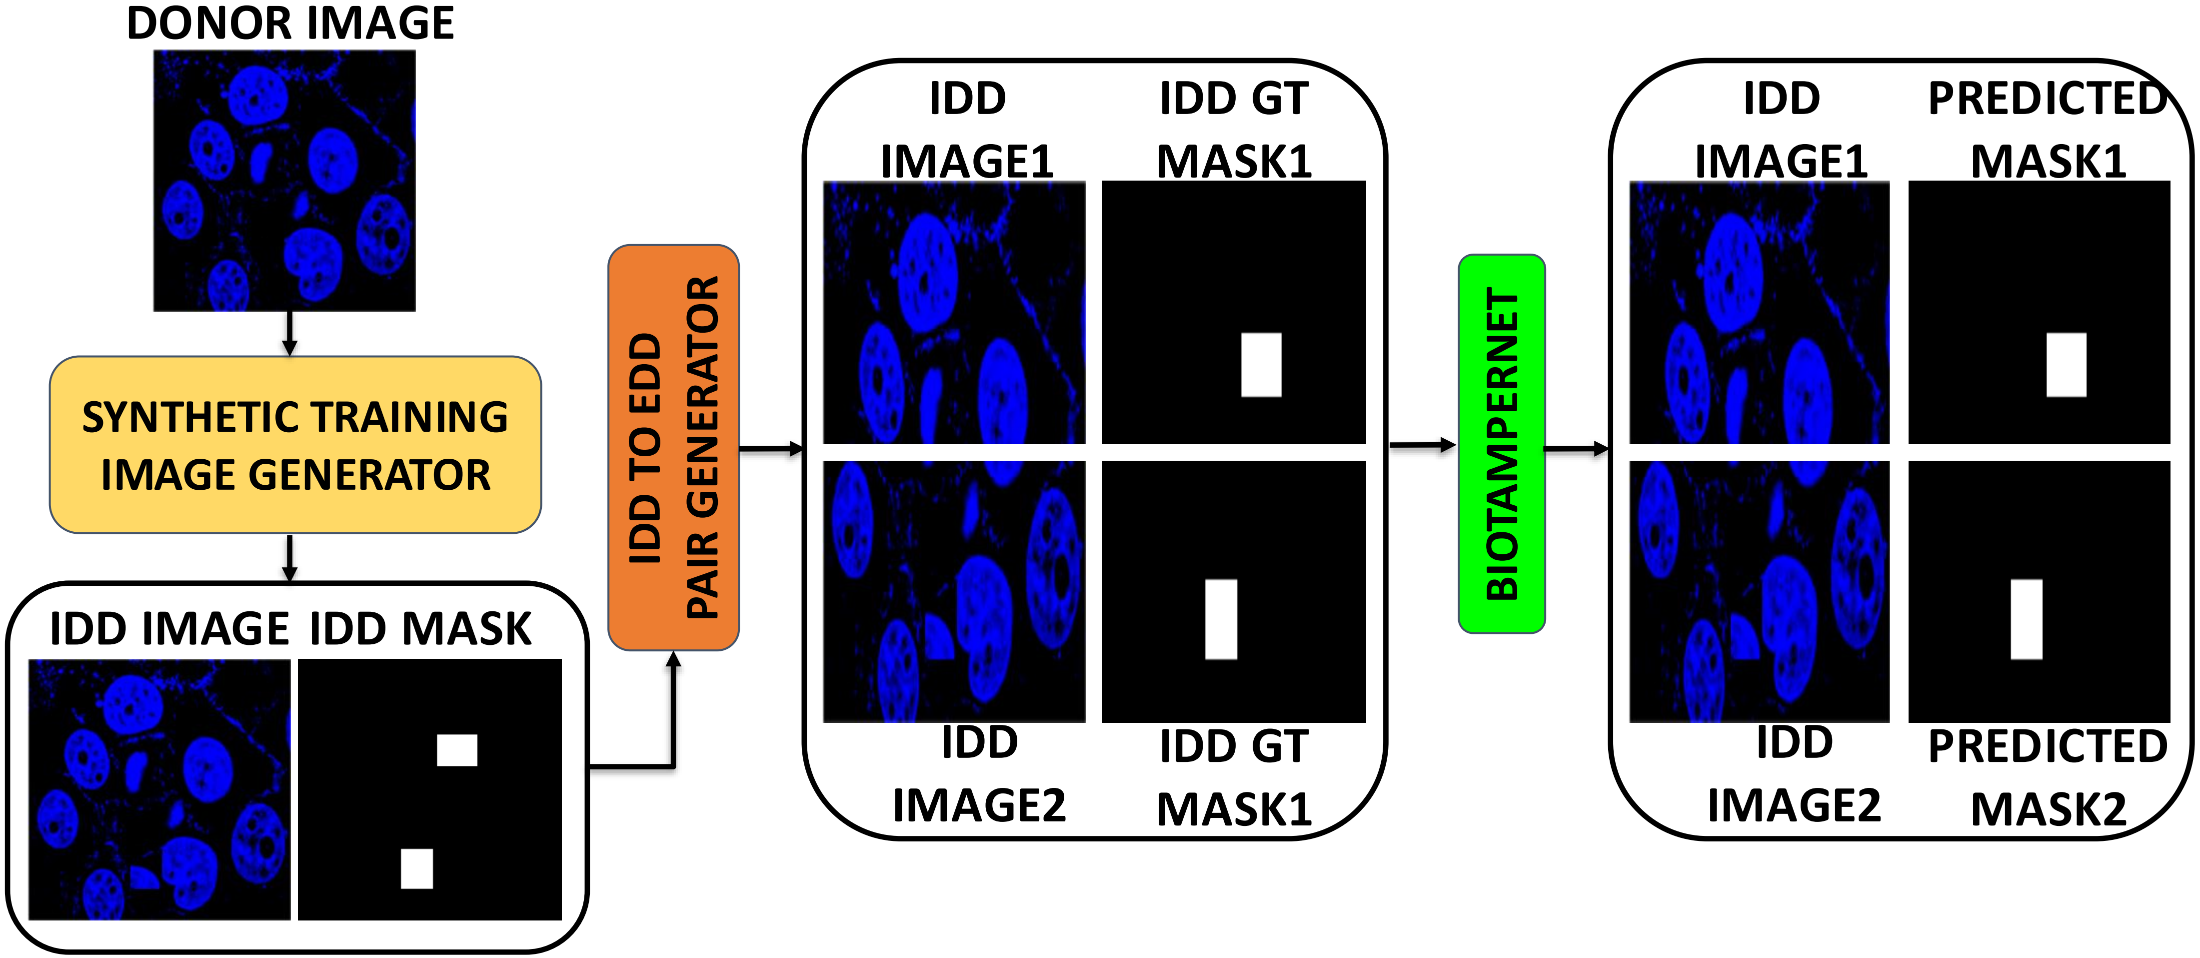

ROADMAP

- What is Scientific Image Misconduct?
    - Bioforensic Image Types
    - Bioforensic Duplication Detection Tasks
    - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
    - BioTamperNet Contributions
    - BioTamperNet Training Setup
    - **State Space Model Overview**
    - BioTamperNet Architecture
    - Experimental Evaluation and Results
    - Directions to Improve

STATE SPACE MODEL OVERVIEW

What is a State Space Model (SSM)?

- A **linear time-invariant (LTI) system** [1] that models sequential dependencies
- Used in signal processing, similar to the **Kalman Filter**
- Defined by **hidden states** evolving over time with input and projection parameters

SSM Formulation:

Given a continuous input *x(t)*, the system computes the output *y(t)* using hidden states *h(t)*:

$$h'(t) = \boldsymbol{A}h(t) + \boldsymbol{B}x(t), \qquad y(t) = \boldsymbol{C}h(t) \ldots\ldots\ldots\ldots (1)$$

A: Governs state evolution
B: Projects input into hidden space
C: Maps hidden states to output

[1] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. ICLR 2022

STATE SPACE MODEL OVERVIEW

Discretization via Zero-Order Hold (ZOH) [1]: Transitions from continuous to discrete SSM

$$\bar{\boldsymbol{A}} = \exp(\Delta \boldsymbol{A}), \quad \bar{\boldsymbol{B}} = (\Delta \boldsymbol{A})^{-1} (\exp(\boldsymbol{A}) - \boldsymbol{I})\, \Delta \boldsymbol{B}, \quad \bar{\boldsymbol{C}} = \boldsymbol{C}$$

$$y_k = \bar{\boldsymbol{C}} h_k + \bar{\boldsymbol{D}} x_k \;, \quad h_k = \bar{\boldsymbol{A}} h_{k-1} + \bar{\boldsymbol{B}} x_k \ldots\ldots\ldots\ldots (2)$$

$\bar{\boldsymbol{D}}$ works as a residual connection and $\bar{\boldsymbol{B}}$ can be approximated using first order Taylor Series as:

$$\bar{\boldsymbol{B}} = (\exp(\boldsymbol{A}) - \boldsymbol{I})\boldsymbol{A}^{-1}\boldsymbol{B} \approx (\Delta \boldsymbol{A})(\Delta \boldsymbol{A})^{-1}\, \Delta \boldsymbol{B} = \Delta \boldsymbol{B}$$

[1] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. ICLR 2022

ROADMAP

- What is Scientific Image Misconduct?
  - ➢ Bioforensic Image Types
  - ➢ Bioforensic Duplication Detection Tasks
  - ➢ Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - ➢ BioTamperNet Contributions
  - ➢ BioTamperNet Training Setup
  - ➢ State Space Model Overview
  - ➢ **BioTamperNet Architecture**
  - ➢ Experimental Evaluation and Results
  - ➢ Directions to Improve

BIOTAMPERNET ARCHITECTURE

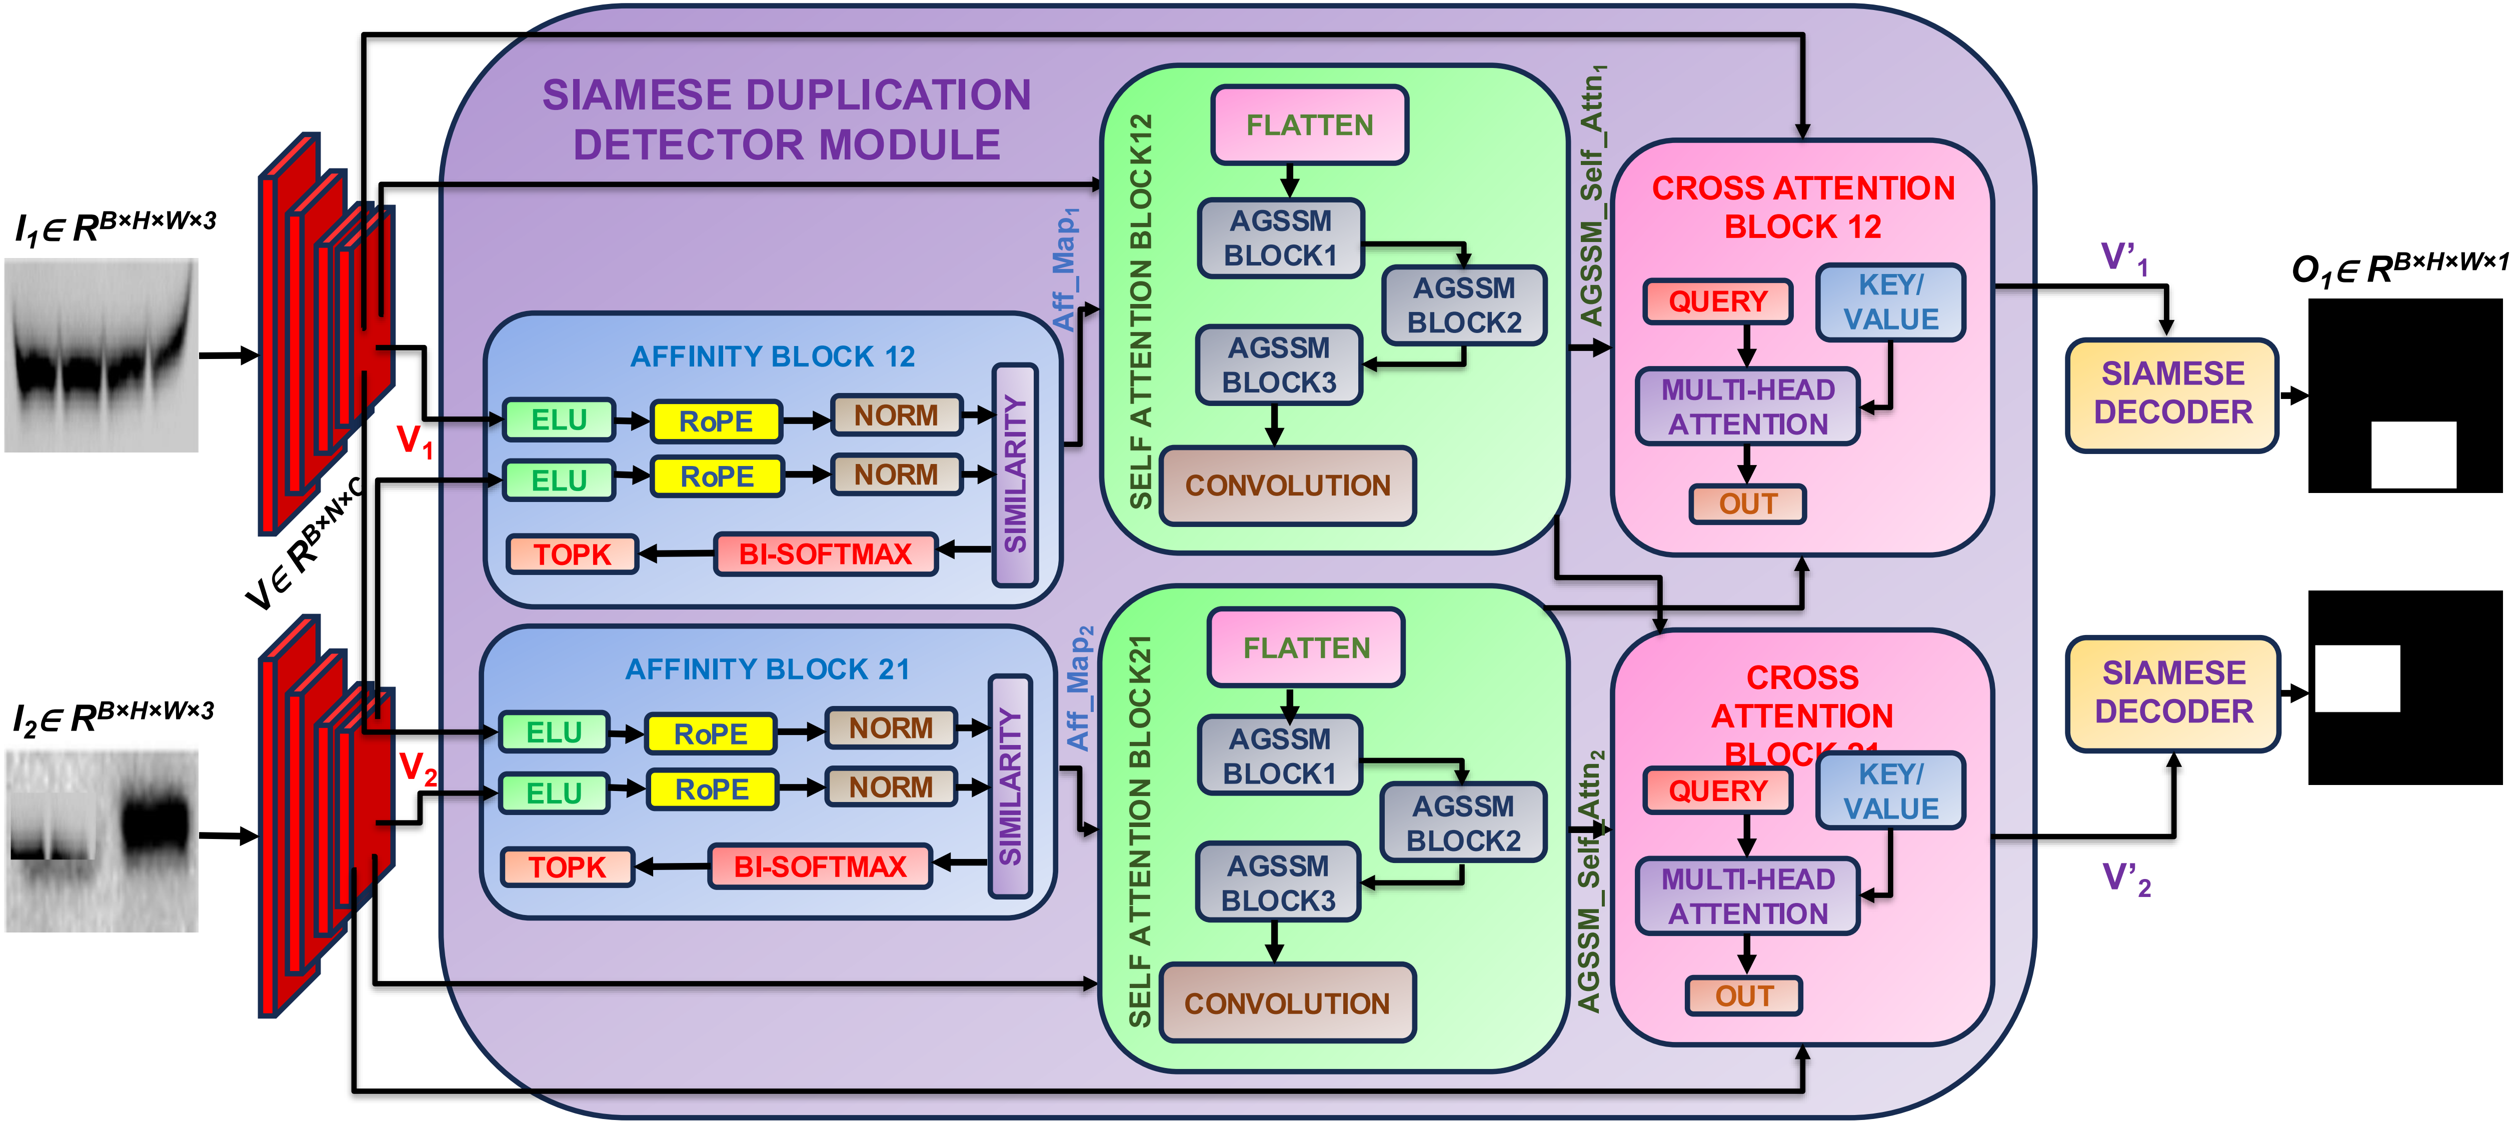

GLOBAL CONTEXT–AWARE STRUCTURED AFFINITY

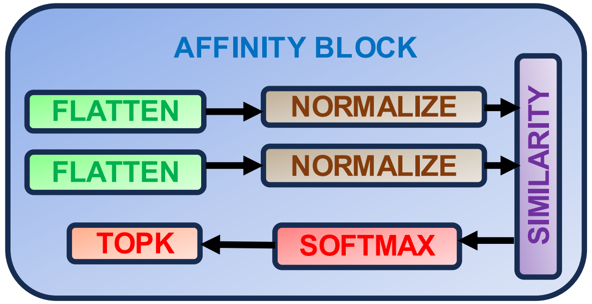

SSM Contextual Encoding

- ➢ Selective-Scan State Space Model (SSM)
- ➢ Aggregates long-range spatial dependencies
- ➢ Context-aware token representations:

$$y_k = \bar{\mathbf{C}} h_k / \bar{\mathbf{C}} n_k + \bar{\mathbf{D}} x_k, \qquad h_k = \bar{\mathbf{A}} h_{k-1} + \bar{\mathbf{B}} x_k \ldots\ldots\ldots\ldots\ldots (2)$$

State-Space Similarity

- ➢ Compute affinity on SSM-encoded features
- ➢ Dot-product similarity:

$$Aff_k = \bar{\mathbf{C}}_k \, \mathbf{B}^T_k \ldots\ldots\ldots (3)$$

- ➢ ELU(V_k) +1.0 for stabilization
- ➢ Rotary Positional Embedding (RoPE) inject spatial inductive bias

Spatial Suppression

- ➢ Mitigates diagonal dominance (self-correlation)
- ➢ Distance-aware kernel modulation:

$$\mathrm{Aff}'_k = \mathrm{Aff}_k \odot K$$

where, $\odot$ is Hadamard Product

Bidirectional Softmax

- ➢ Row-wise and Column-wise normalizations
- ➢ Enforces mutual similarity consistency

Convolutional Refinement + Top-K

- ➢ 4-layer convolutional refinement
- ➢ Structured affinity map
- ➢ Retain strongest duplication cues

AFFINITY-GUIDED SELF-ATTENTION

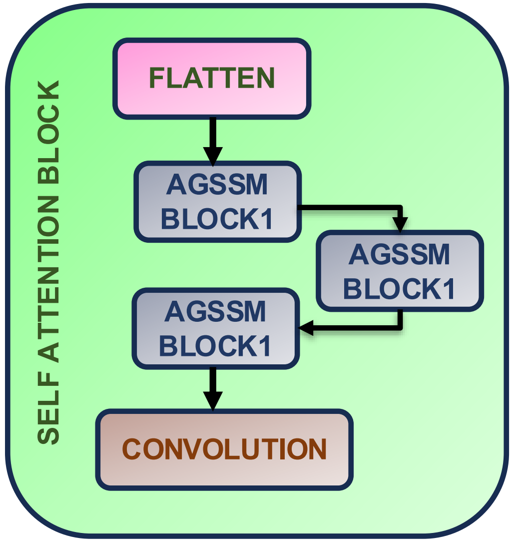

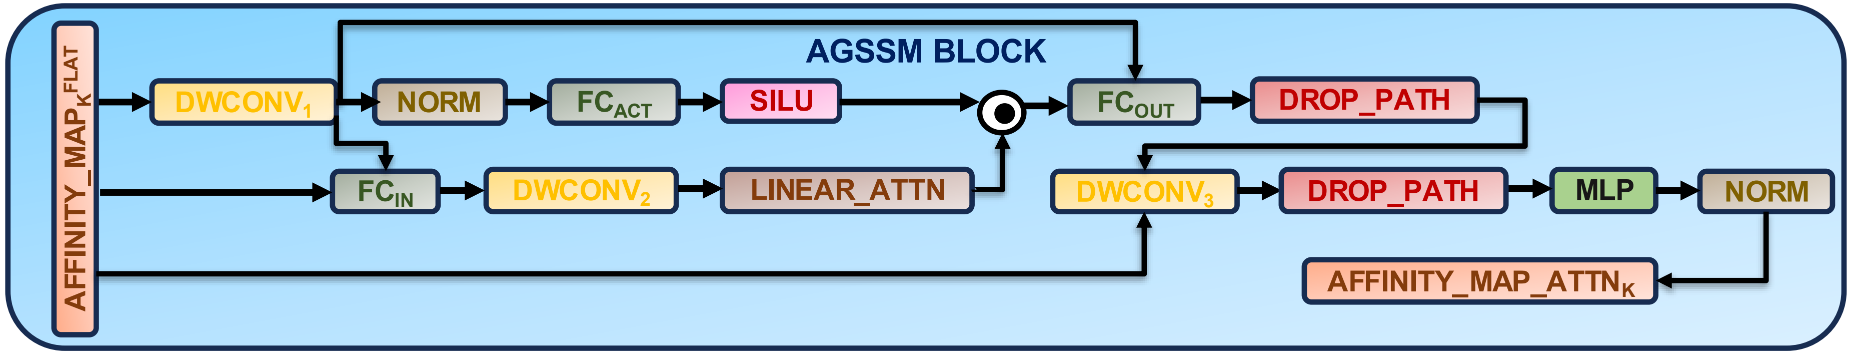

AGSSM Block (Per Branch)

- SSM-based Linear Attention
- Dynamic Gating via:

  $a = SiLU(FC(Norm(Affinity)))$
- Residual + DropPath for stabilization

Affinity-Modulated Attention

- Flattened affinity map $Affinity\_Map_{flat} \in \mathbb{R}^N$
- Injected directly into attention computation
- Modulates token interactions using duplication priors

Parallel AGSSM Modules

- 3 parallel AGSSM blocks learn complementary duplication-aware interaction patterns
- Promote diverse local interactions across receptive fields and gating patterns
- Averaged aggregation Stabilizes training and reduces sensitivity to noisy affinity cues:

$$1/3 \sum_{i=1}^{3} AGSSM_i(Affinity\_Map\_Attn_k)$$

Final Projection

- 1×1 convolution mixes channel information and aligns feature dimensions
- Residual fusion preserves original semantics while injecting affinity-guided refinement:

$$V'_k = Vk + AGSSM\_Self\_Attn_k$$

AFFINITY-GUIDED CROSS-ATTENTION

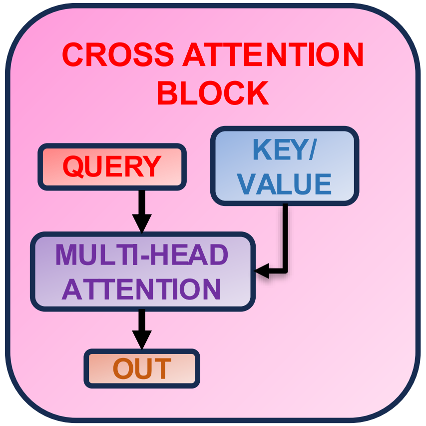

CROSS ATTENTION BLOCK

QUERY

KEY/ VALUE

MULTI-HEAD ATTENTION

OUT

Cross-View Feature Interaction

- Cross-attention propagates duplication-aware cues across image pairs
- Promotes information exchange between paired inputs
- Aligns duplicated regions across views
- Operates after affinity-guided self-attention refinement

Query–Key Formulation

- Q = Flatten($AGSSM\_Self\_Attn_1$)
- K = Flatten($Conv_{1\times 1}$ $(AGSSM\_Self\_Attn_2$))
- Multi-head cross-attention:

$$CrossAttn(Q,K,K)$$

Projected Key Alignment

- 1×1 convolution aligns feature dimensions
- Encourages structured cross-view matching

Residual Cross-View Fusion

$$V1' = V1 + \mathrm{CrossAttn}(V1,V2)$$
$$V2' = V2 + \mathrm{CrossAttn}(V2,V1)$$

- Symmetric refinement
- Enhances contextual consistency

ROADMAP

- What is Scientific Image Misconduct?
  - Bioforensic Image Types
  - Bioforensic Duplication Detection Tasks
  - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - BioTamperNet Contributions
  - BioTamperNet Training Setup
  - State Space Model Overview
  - BioTamperNet Architecture
  - **Experimental Evaluation and Results**
  - Directions to Improve

QUALITATIVE EVALUATION - EDD

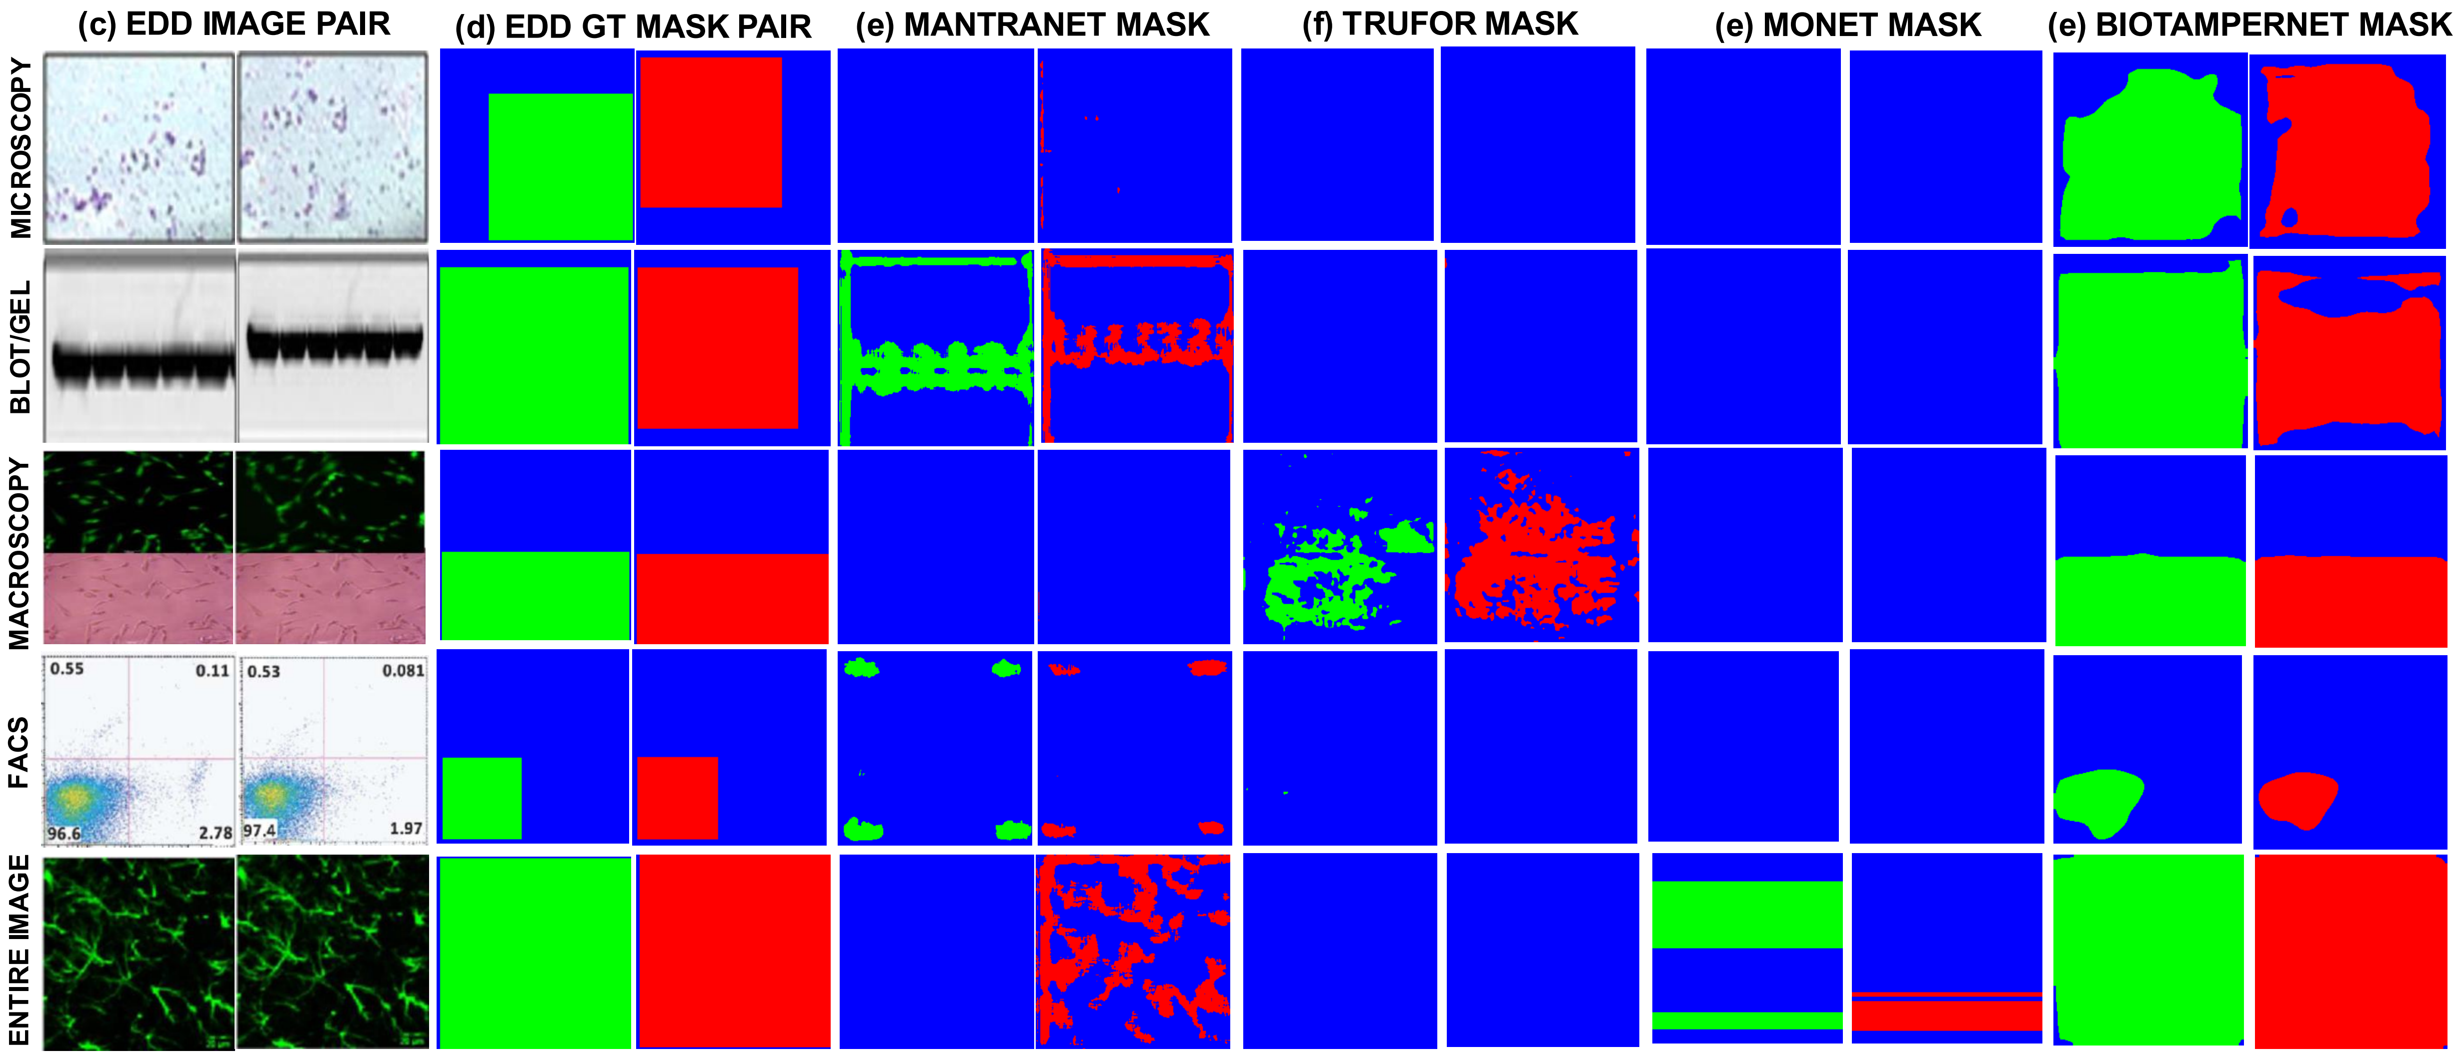

QUALITATIVE EVALUATION - IDD

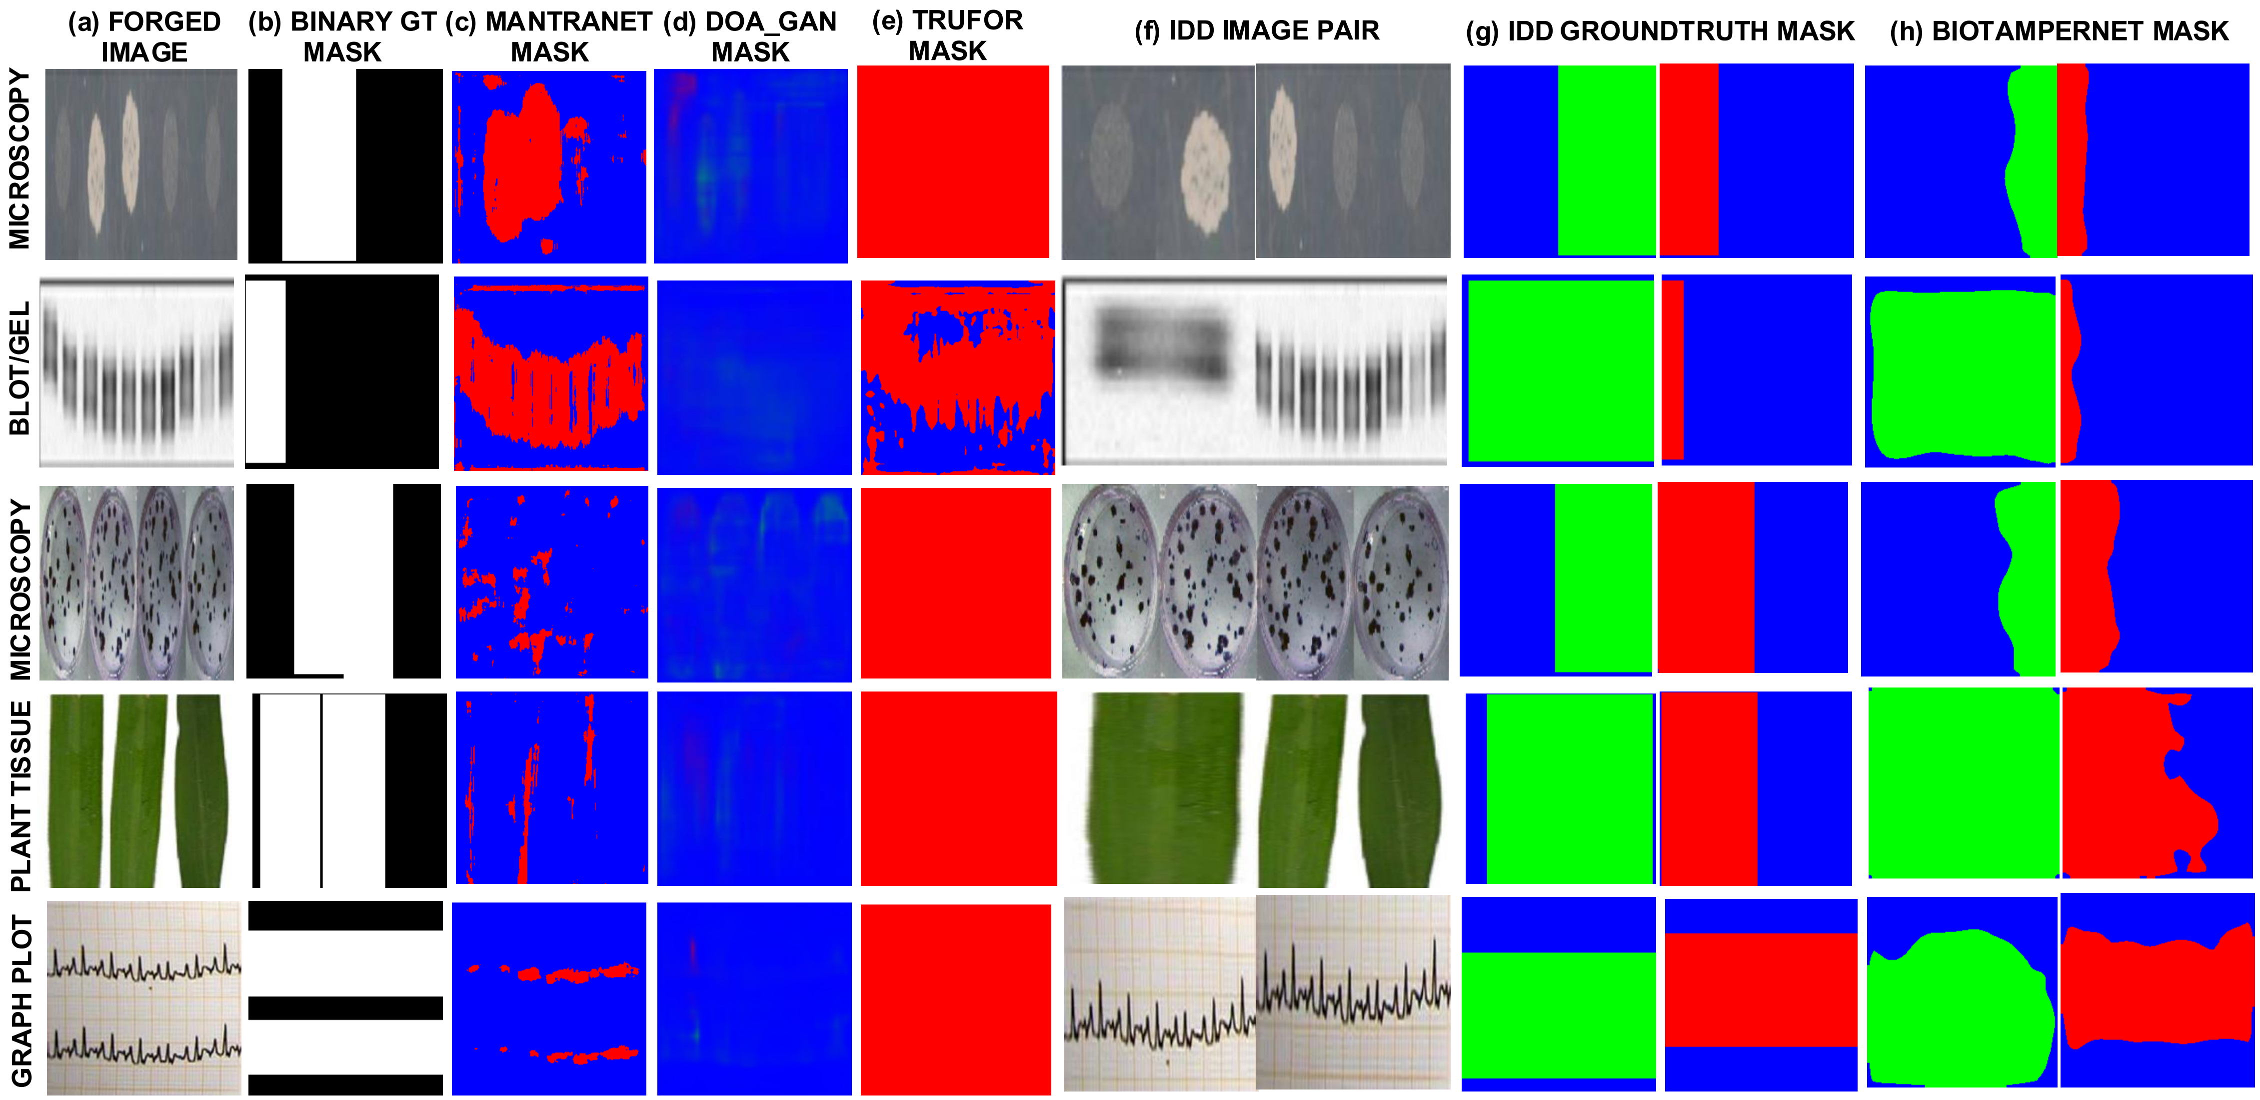

BIOTAMPERNET AFFINITY HEATMAPS

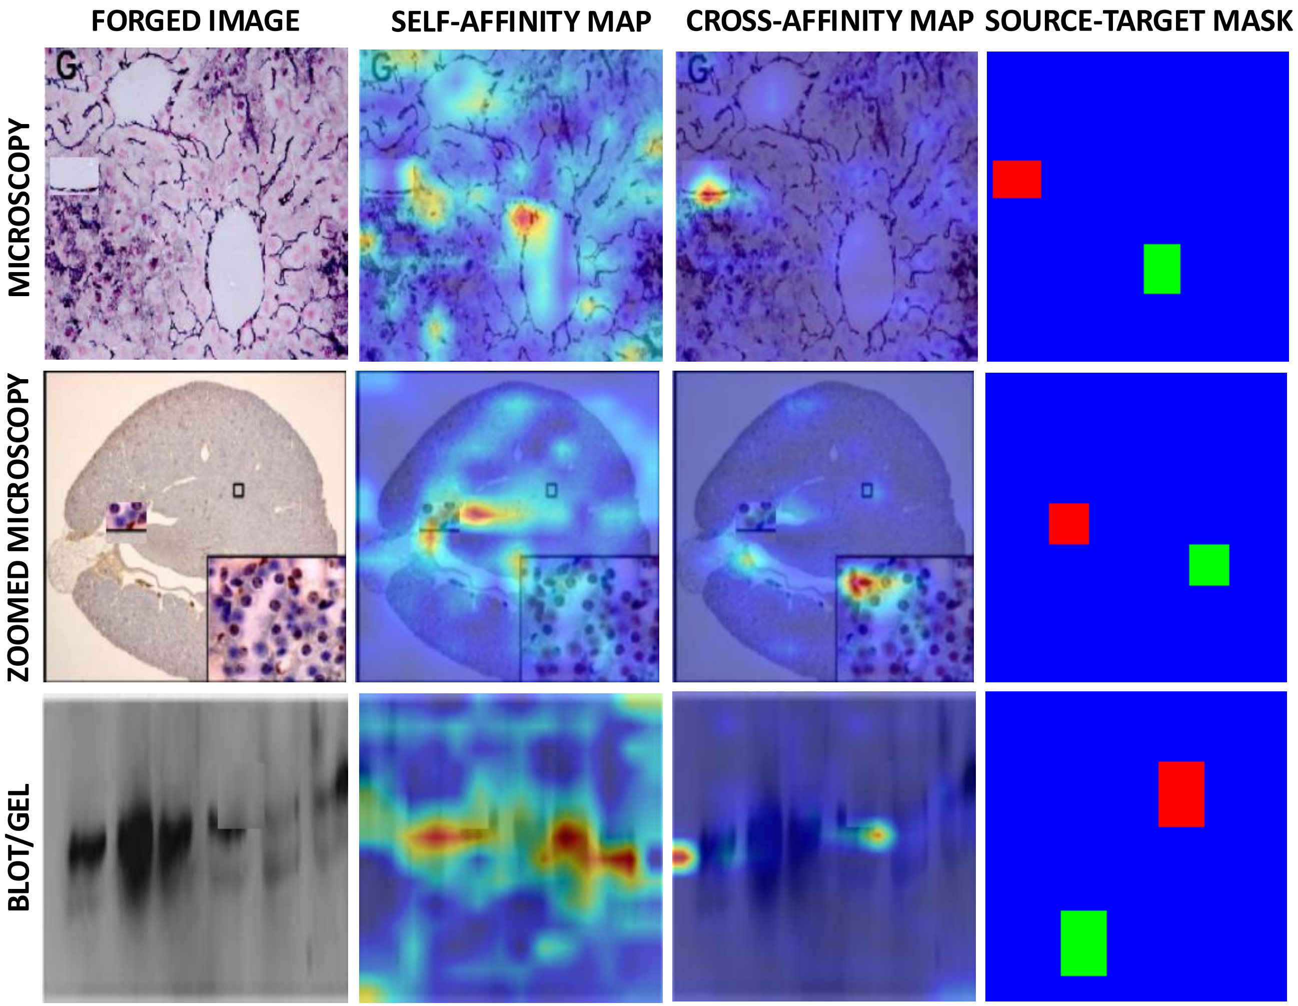

- **Self-affinity** highlights repeated structures within the same image
- **Cross-affinity** links duplicated source and target regions
- High-response areas align closely with ground-truth masks
- Consistent behavior across multiple biomedical modalities
- Demonstrates meaningful structural similarity learning for duplication localization

Fig. 1: Interpretable Affinity Heatmaps for Duplication Localization - Ground-truth Untampered, Source, and Target regions

BIOTAMPERNET VS BASELINES - EDD

Method	Conference	Microscopy		Blot/gel		Macroscopy		Facs		Combined	
		Image	Pixel	Image	Pixel	Image	Pixel	Image	Pixel	Image	Pixel
SIFT	IJCV 2004	0.180	0.146	0.113	0.148	0.130	0.194	0.110	0.073	0.142	0.132
DenseField-ZM	TIFS 2015	0.422	0.425	0.161	0.192	0.285	0.256	0.540	**0.504**	0.278	0.324
DMVN	ACM MM 2017	0.342	0.242	0.430	0.261	0.238	0.185	0.282	0.164	0.310	0.244
ManTraNet	CVPR 2019	0.347	0.244	0.449	0.287	0.275	0.202	0.337	0.186	0.351	0.231
TruFor	CVPR 2023	0.381	0.264	0.485	0.311	0.257	0.358	0.363	0.194	0.371	0.282
MONet	ICIP 2022	0.435	0.398	0.520	0.507	0.262	0.221	0.356	0.313	0.438	0.410
SparseViT	AAAI 2025	0.384	0.269	0.482	0.309	0.271	0.371	0.376	0.202	0.378	0.288
BioTamperNet	ICLR 2026	**0.739**	**0.487**	**0.672**	**0.589**	**0.743**	**0.577**	**0.652**	0.448	**0.701**	**0.526**

Table 2: Image and Pixel Level MCC Scores on External Duplication Detection (EDD) in BioFors Test set from retracted publications. All models trained on BioFors train set. ***Bold - Best****, Underline - Second Best*

BIOTAMPERNET VS BASELINES - IDD

Method	Conference	Microscopy		Blot/gel		Macroscopy		Combined	
		Image	Pixel	Image	Pixel	Image	Pixel	Image	Pixel
DenseField-ZM	TIFS 2015	0.764	0.197	0.515	0.449	0.573	0.478	0.564	0.353
DenseField-PCT	TIFS 2015	0.638	0.167	0.503	0.466	0.712	0.487	0.569	0.364
DenseField-FMT	TIFS 2015	0.183	0.178	0.480	0.400	0.495	0.458	0.509	0.316
BusterNet	ECCV 2018	0.183	0.178	0.226	0.076	0.021	0.106	0.269	0.107
ManTraNet	CVPR 2019	0.316	0.194	0.317	0.094	0.272	0.262	0.335	0.183
TruFor	CVPR 2023	0.336	0.199	0.351	0.106	0.324	0.297	0.337	0.201
SparseViT	AAAI 2025	0.342	0.211	0.356	0.117	0.331	0.305	0.343	0.211
BioTamperNet	ICLR 2026	**0.827**	**0.526**	**0.681**	**0.617**	**0.843**	**0.679**	**0.701**	**0.534**

Table 3: Image and Pixel Level MCC Scores on Internal Duplication Detection (IDD) in BioFors Test set from retracted publications. All models trained on BioFors train set. ***Bold - Best****,* *Underline - Second Best*

BIOTAMPERNET VS BASELINES - CSTD

Method	Conference	F1		MCC	
		Image	Pixel	Image	Pixel
ManTraNet	CVPR 2019	0.253	0.090	0.170	0.080
TruFor	CVPR 2023	<u>0.311</u>	<u>0.104</u>	<u>0.173</u>	<u>0.092</u>
BioTamperNet	ICLR 2026	**0.537**	**0.378**	**0.514**	**0.346**

Table 3: CSTD Detection on BioFors Test Set. All models trained on BioFors train set.
***Bold - Best**, <u>Underline - Second Best</u>*

SCIENTIFIC INTEGRITY IN SYNTHETIC BENCHMARKS

Method	Conference	RSIID [Carduento & Rocha, Science and Engineering Ethics, 2022]	Western Blots [Manjunath et al. IEEE WIFS 2024]	Average
ManTraNet	CVPR 2019	0.462	0.196	0.284
TruFor	CVPR 2023	0.825	0.457	0.553
SparseViT	AAAI 2025	<u>0.842</u>	<u>0.739</u>	<u>0.612</u>
BioTamperNet	ICLR 2026	**0.965**	**0.913**	**0.628**

Table 3: Pixel-Level MCC Scores on Forgery Detection for Generative AI based Synthetic Scientific Integrity Datasets. All models trained on BioFors train set. ***Bold - Best****,* *<u>Underline - Second Best</u>*

BIOTAMPERNET ABLATION STUDIES

Method	Microscopy	Blot/Gel	Macroscopy
w/o Affinity	0.421	0.489	0.462
w/o SSM (CNN)	0.393	0.453	0.437
w/o SSM (ViT-MHA)	0.407	0.466	0.445
w/o Self-Attn	0.451	0.509	0.492
w/o Cross-Attn	0.444	0.497	0.481
w/ LayerNorm	0.439	0.492	0.476
BioTamperNet (Full)	**0.487**	**0.589**	0.577
BioTamperNet (Full) + Global SSM	0.467	0.539	**0.580**

Table 5: MCC scores using BioTamperNet (trained on BioFors EDD Train set) and evaluated on BioFors EDD Test set. ***Bold - Best****, Underline - Second Best*

BIOTAMPERNET ROBUSTNESS ANALYSIS

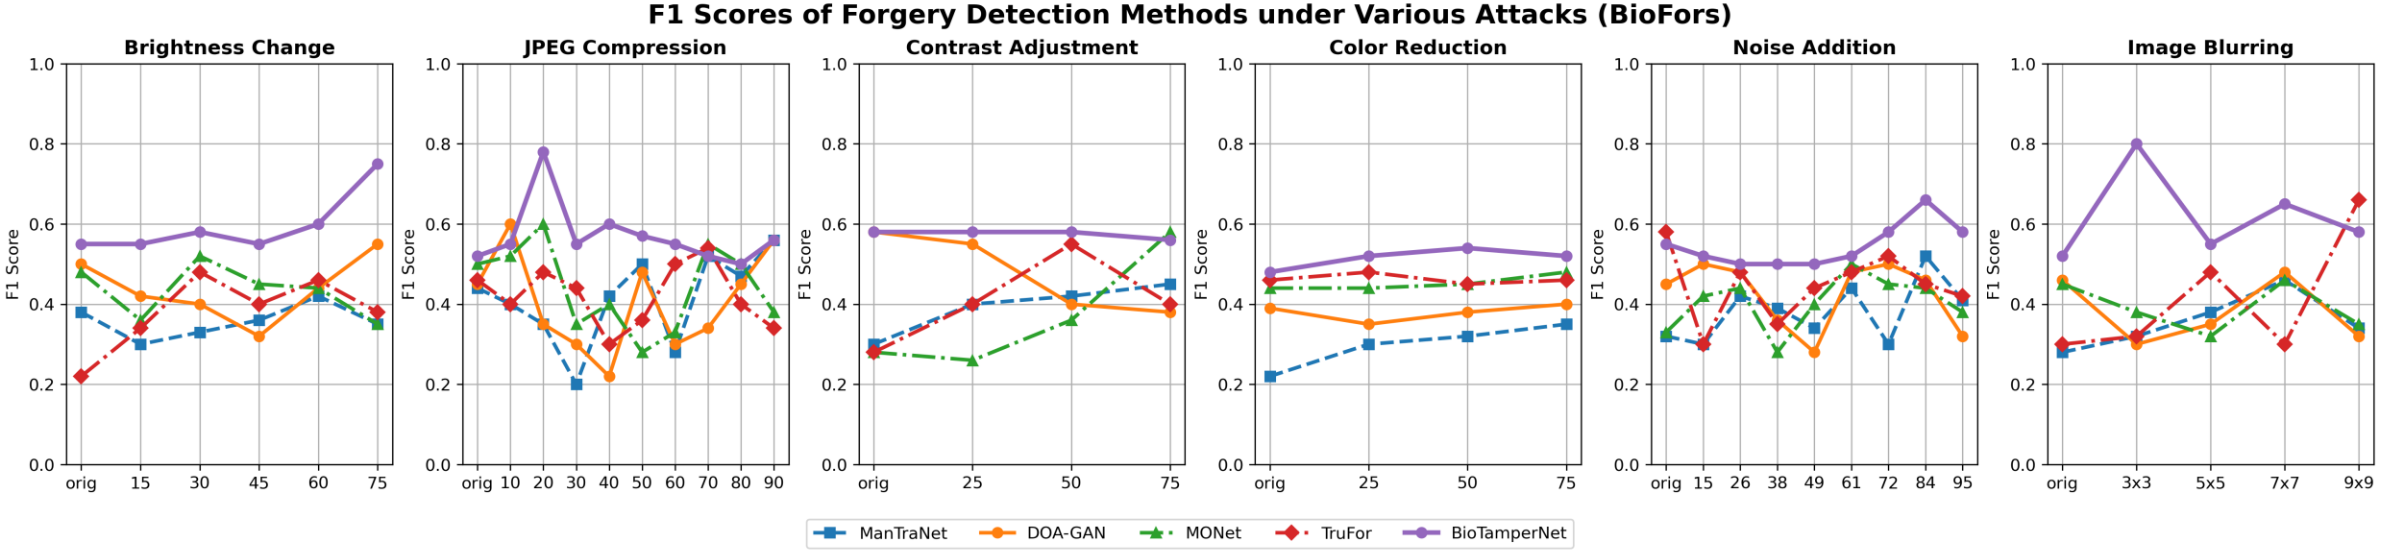

Fig 13: Robustness analysis of BioTamperNet compared to baselines on the Biofors [ICCV 2021] dataset across six common perturbations—following the protocol in ManTraNet [CVPR 2019]

BIOTAMPERNET MODEL COMPLEXITY

Method	Conference	Backbone	Input Size	Parameters	FLOPs
BusterNet	ECCV 2018	VGG	256 X 256	15.5 M	45.7 G
ManTraNet	CVPR 2019	VGG	256 X 256	3.9 M	274.0 G
MONet	ICIP 2022	ResNet	512 X 512	31.0 M	78.0 G
TruFor	CVPR 2023	ViT	512 X 512	68.7 M	236.5 G
SparseViT	AAAI 2025	ViT	512 X 512	50.3 M	46.2 G
BioTamperNet	ICLR 2026	SSM	512 X 512	**36.7 M**	**29.6 G**

Table 7: Model Complexity with respect to Resolution, Parameters, FLOPS.
Bold - Best, *Underline - Second Best*

ROADMAP

- What is Scientific Image Misconduct?
  - Bioforensic Image Types
  - Bioforensic Duplication Detection Tasks
  - Why Biomedical Forgery Detection is Hard?
- Solution Proposed – BioTamperNet Model
  - BioTamperNet Contributions
  - BioTamperNet Training Setup
  - State Space Model Overview
  - BioTamperNet Architecture
  - Experimental Evaluation and Results
  - **Directions to Improve**

BIOTAMPERNET LIMITATIONS

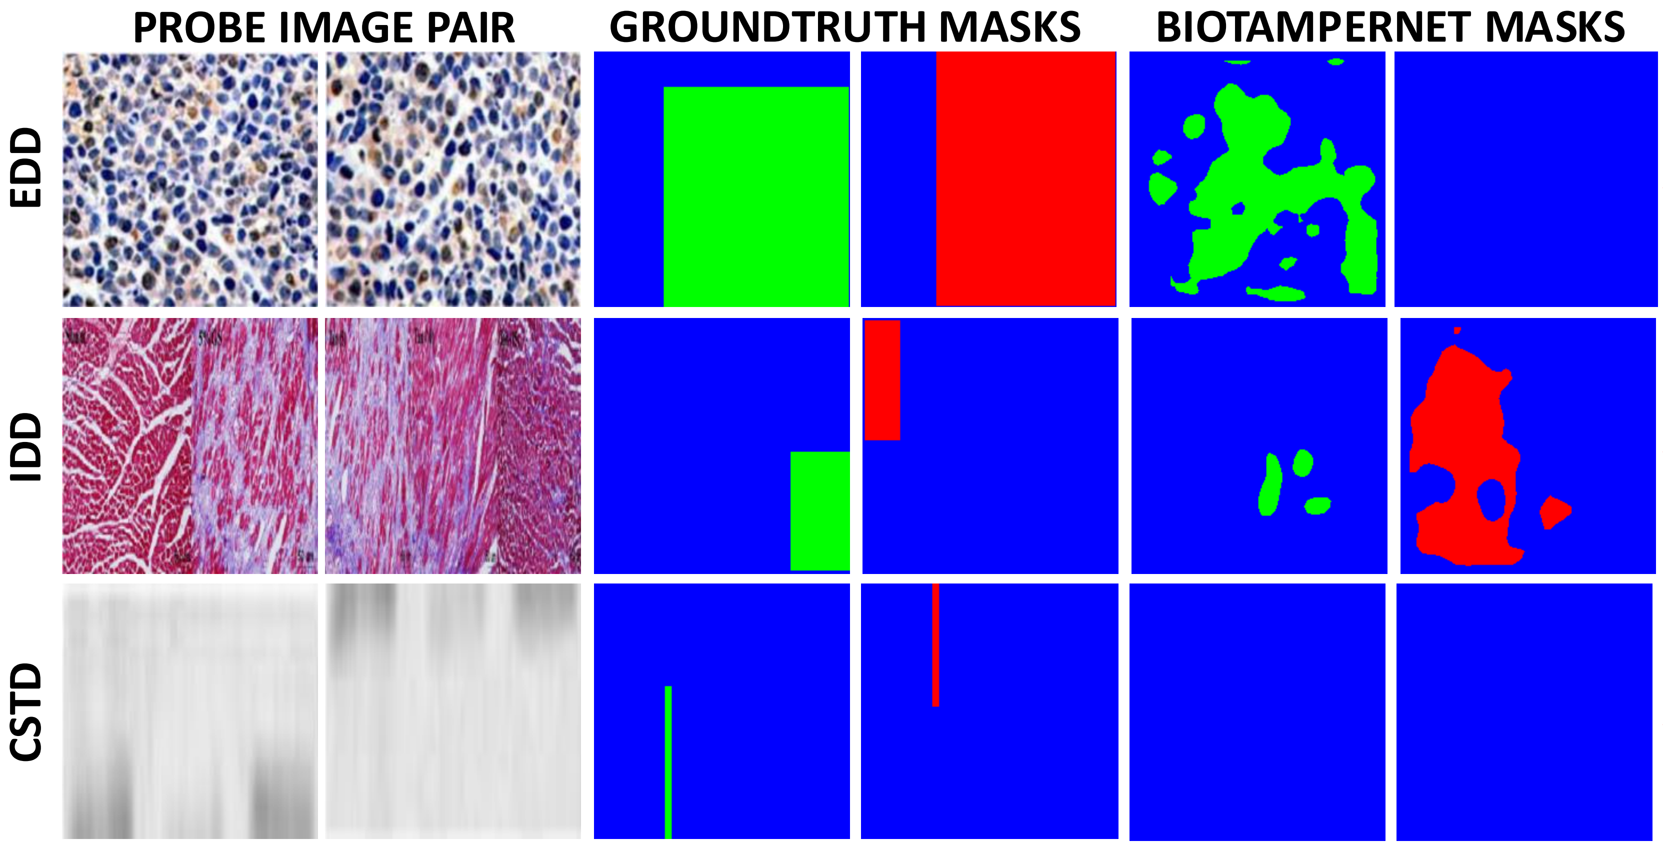

Overlapping source–target regions
→ boundary ambiguity

False positives in repetitive
textures / clustered structures

Blot/Gel images with dense stains
→ weak boundary contrast

DIRECTIONS TO IMPROVE

Planned Improvements

- Overlap-aware post-processing + auxiliary boundary heads
- Semantic regularization (cycle-consistency, entropy mining, adaptive thresholding)
- Frequency-aware filtering, 2D RoPE, localized NMS
- Stain-invariant augmentation + multi-scale refinement

Future Direction

- Extend to video-level forgery detection with temporal attention and spatio-temporal consistency

Questions?

Thank You