

# **Bio-Inspired Identity Representation Learning via Dual-Attention Feature Decoupling and Cellular Phase-Separation Dynamics for Robust Visual Re-Identification**

Rowan L. Makinen

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.  
rowanwork@uc.edu

Castian Wolfe

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.  
wolfe473@uab.edu

Kexiao Long

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.  
kexiaomail@buffalo.edu

Walid Graham

Department of Computer Science, Binghamton University, Binghamton, NY, USA.  
grahamwalid@binghamton.edu

## **Abstract**

Visual re-identification faces persistent challenges arising from clothing changes, viewpoint variations, and adversarial occlusions, demanding representation learning paradigms that move beyond rigid feature extraction toward dynamic, context-aware identity encoding. This work proposes a bio-inspired identity representation framework that synergizes dual-attention feature decoupling with cellular phase-separation dynamics to achieve robust, clothing-invariant person re-identification. The system-level architecture decouples identity-discriminative cues from environmental and appearance confounders through parallel spatial and channel attention streams, while a phase-separation-inspired dynamic reconfiguration module organizes the learned feature manifold into condensed, identity-consistent clusters that resist feature drift across temporal and contextual gaps. Rather than focusing on algorithmic minutiae, the paper presents a comprehensive analysis of structural trade-offs, governance, deployment sustainability, fairness, and socio-technical implications inherent in large-scale re-identification infrastructures. The discussion examines how biomimetic self-organization principles can promote resilience against distribution shifts and adversarial manipulations, how modular decoupling enables transparent auditing and bias mitigation, and how edge-cloud orchestration constrains energy footprints while preserving real-time throughput. The paper further reflects on the broader policy landscape surrounding biometric surveillance, emphasizing the dual-use dilemma and the urgent need for embedded fairness-by-design and continuous regulatory alignment in identity-aware systems.

## **Keywords**

visual re-identification, bio-inspired computing, dual-attention decoupling, phase separation, robustness, fairness, system architecture, governance, sustainable AI.

## 1. Introduction

Visual person re-identification has matured into a critical enabler of smart surveillance, autonomous retail, and urban analytics, yet its underlying representation learning remains fragile under the combinatorial complexity of real-world appearance variations. Clothing alterations, illumination shifts, partial occlusions, and viewpoint changes systematically degrade the performance of models trained under closed-world assumptions, exposing fundamental limitations in static feature encoding schemes. Recent advances in attention-based architectures and feature disentanglement have pushed accuracy boundaries on curated benchmarks, but the systemic robustness, fairness, and deployability of such models in large-scale, open-set environments remain under-investigated from a systems perspective. There is an urgent need to reimagine identity representation not as a fixed embedding but as an adaptive, self-organizing process that can maintain consistency across time while filtering out non-identity confounders. Biology provides an instructive model: cellular systems achieve remarkable molecular identity recognition and signal filtering through phase-separated condensates that dynamically organize proteins and nucleic acids without central coordination. Drawing on this principle, the present paper proposes a bio-inspired identity representation framework that integrates dual-attention feature decoupling with cellular phase-separation dynamics at the system level, aiming to deliver robust, clothing-agnostic re-identification under adversarial and diverse deployment conditions.

The contributions of this work are positioned at the intersection of system design, architectural governance, and socio-technical assessment rather than at purely algorithmic optimization. Through conceptual analysis and cross-domain synthesis, the paper elucidates how decoupling identity-relevant signals from clothing and environmental biases can be structurally implemented using parallel attention streams, and how phase-separation-inspired dynamic clustering can stabilize the feature manifold against perturbations commonly encountered in longitudinal person tracking. The exposition foregrounds trade-offs among computational cost, interpretability, modularity, and fairness, and it connects these architectural choices to larger deployment ecosystems involving edge computing, regulatory compliance, and sustainability mandates. By framing the system as a socio-technical infrastructure, the discussion extends to the ethical implications of identity-aware computer vision, proposing governance mechanisms that leverage architectural transparency for auditing and de-biasing. The paper thus serves as a systems blueprint for next-generation re-identification platforms that are simultaneously robust, accountable, and ecologically viable.

## 2. Related Work and Theoretical Foundations

Person re-identification research has progressed from handcrafted feature pipelines to deep learning baselines that emphasize metric learning and part-based features [1, 2]. The introduction of attention mechanisms, originally popularized in natural language processing, catalyzed a shift toward context-aware feature recalibration in vision systems [3]. Subsequent works exploited spatial and channel attention to suppress background clutter and highlight discriminative body parts, establishing dual-attention as a powerful tool for fine-grained recognition. In parallel, feature decoupling emerged as a strategy to separate identity-preserving information from nuisance variability [4]. By forcing the network to explicitly model and discard clothing, pose, and illumination factors, decoupling reduces overfitting to spurious correlations and improves cross-domain generalization. Despite these advances,

most approaches treat decoupling as a static architectural constraint, failing to account for the dynamic, non-stationary nature of identity signals over time and space.

Cellular biology offers a compelling conceptual reservoir through the phenomenon of liquid-liquid phase separation, where multivalent interactions drive the formation of membraneless organelles that spatially organize biochemical reactions [5]. Recent work has shown that phase separation of transcriptional co-activators such as YAP-MAML2 differentially regulates gene expression programs, exemplifying how physical condensation can dynamically route information flow in complex environments. The notion that functional specificity can arise from the collective dynamics of weakly interacting components without a centralized controller resonates with the challenge of organizing person representations across disjoint camera views. Translating this principle into a computational framework suggests that identity features could be actively reconfigured through affinity-driven condensation processes that segregate identity-consistent points from noise-induced outliers, thereby improving robustness to clothing changes and occlusions.

Within the machine learning community, the idea of dynamic feature organization has been explored through attention routing, capsule networks, and graph neural diffusion, yet a systematic integration of biological condensation dynamics with dual-attention decoupling for re-identification has remained absent. Prior art in clothing-change re-identification has recognized the necessity of multimodal fusion and feature-driven decoupling, as demonstrated by frameworks that combine biometric shape cues with texture-invariant representations [8]. Meanwhile, dual-attention architectures have been refined in other high-stakes image analysis domains, including medical image segmentation [9]. Still, the translation of such attention designs into a robust, deployable socio-technical system for re-identification requires careful consideration of infrastructure, fairness, and governance layers that extend far beyond component-level performance.

### **3. System Architecture: Dual-Attention Feature Decoupling and Cellular Phase-Separation Dynamics**

The proposed system architecture is organized around two interdependent subsystems that jointly realize robust identity encoding: a dual-attention feature decoupling front-end and a phase-separation-inspired dynamic condensation back-end. The dual-attention module comprises parallel spatial and channel attention pathways that operate on intermediate convolutional feature maps. The spatial attention stream learns a two-dimensional saliency map that amplifies body regions likely to convey stable identity cues, such as skeletal proportions and gait-related limb configurations, while suppressing transient regions dominated by clothing texture and accessories. Concurrently, the channel attention stream recalibrates feature channel importance by capturing long-range interdependencies among filter responses, effectively re-weighting the representational budget toward identity-discriminative channels that are invariant to common appearance perturbations. The outputs of both streams are fused and subsequently decomposed into identity-associated and clothing-associated embedding subspaces through a gating mechanism that enforces informational bottlenecks on the clothing branch. This decoupling is not simply a softened separation loss but an architectural commitment that can be audited and probed at inference time, an essential property for accountable systems.

Following decoupling, the identity-associated embeddings enter the cellular phase-separation dynamics module, which is designed to mimic the formation and dissolution of biomolecular condensates. Rather than employing a fixed distance metric or static clustering, this module

treats each embedding as a particle capable of forming multivalent interactions with its neighbors in a high-dimensional similarity space. An iterative process, governed by interaction strengths derived from learned affinity kernels and density-dependent refinement, causes embeddings that share genuine identity consistency to progressively condense into compact clusters while repelling noise points and cross-identity interference. The dynamics are designed to be locally computed and globally coherent, allowing the system to update identity representations online as new gallery observations arrive without full retraining. This self-organizing behavior draws inspiration from how phase-separated YAP-MAML2 condensates modulate transcriptional output, achieving selective signal amplification through physical association rather than explicit labeling.

From a systems standpoint, the dual-attention decoupling and phase-separation dynamics introduce several structural trade-offs. The decoupling front-end increases parameter count and memory footprint relative to a monolithic backbone, but the modular decomposition facilitates per-component optimization for distinct hardware accelerators, such as running spatial attention on dedicated edge neural processing units while channel attention is handled in the cloud. The condensation back-end requires maintaining a dynamic neighborhood graph, whose computational complexity scales quadratically with gallery size in naive implementations, yet approximation techniques including locality-sensitive hashing and online clustering reduce the burden to near-linear time. Architects must balance the refresh frequency of the condensation process against latency constraints: higher refresh rates improve resilience to distribution drift but consume more energy, posing a direct tension between robustness and sustainability. Moreover, the decoupled clothing branch can be retained or discarded based on institutional privacy policies; in high-surveillance contexts, its removal serves as a privacy-preserving measure, while in fashion analytics it can be repurposed for secondary inferences. This architectural flexibility exemplifies how bio-inspired design can be tailored to compliance requirements through structural modularity.

#### **4. Governance, Robustness, and Fairness Implications**

Adopting a bio-inspired, dynamically self-organizing identity representation carries profound implications for system governance, particularly regarding robustness and fairness. Robustness to adversarial perturbations and domain shifts is enhanced by two orthogonal mechanisms: the attention-driven decoupling that suppresses spurious correlations, and the condensation-based manifold stabilization that prevents isolated erroneous embeddings from propagating identity errors across camera networks. When an adversary introduces carefully crafted patches or lighting artifacts intended to mimic a target identity, the spatial attention mechanism can down-weight manipulated regions if they deviate from learned body-structural priors, while the phase-separation module tends to exclude such outliers from identity-consistent condensates because their affinity profiles differ from those of authentic embeddings. This dual guard offers a defense-in-depth strategy that is structurally distinct from adversarial training alone, which often remains vulnerable to unforeseen attack surfaces.

Fairness considerations permeate the entire pipeline. Person re-identification systems are known to exhibit performance disparities across demographic groups, partly because training data imbalances correlate with clothing styles, body shapes, and co-occurrence biases. The decoupling architecture creates an explicit separation between identity cues and clothing-related features, which can be leveraged for fairness auditing. By monitoring the activation patterns in the clothing branch, system auditors can detect whether the model inadvertently associates certain communities with specific garment categories in a discriminatory manner.

Furthermore, the phase-separation dynamics can be tuned to enforce equitable cluster condensation by adaptively adjusting affinity thresholds in a group-sensitive fashion, ensuring that marginalized identity distributions do not remain fragmented simply because they are underrepresented. This operationalizes a form of procedural fairness embedded in the system's self-organization logic, shifting the burden of bias mitigation from post-hoc recalibration to online dynamic governance. However, such group-sensitive mechanisms must themselves be carefully designed to avoid encoding protected attributes explicitly, a challenge that demands ongoing collaboration between system architects, ethicists, and domain regulators.

The broader policy landscape surrounding re-identification technologies shapes deployment decisions and governance frameworks. In jurisdictions with strict biometric privacy laws, the ability to discard the clothing feature branch or to audit the decoupling process can serve as a verifiable technical safeguard that eases regulatory approval. Conversely, the condensation module's capacity to maintain persistent identity profiles across time and space intensifies surveillance potential, raising the specter of misuse by authoritarian regimes. System designers must therefore embed mandatory access controls and lineage-tracking logs that constrain how condensed identity clusters are queried and retained. International technology transfer agreements and certification regimes such as ISO/IEC 42001 for AI management systems could reference architectural transparency metrics derived from the decoupling and phase-separation design, thus bridging technical innovation and global governance.

## **5. Deployment and Sustainability**

Deploying a computationally intensive bio-inspired re-identification system at urban scale demands a thoughtful edge-cloud orchestration that reconciles throughput, latency, and energy consumption. Camera-side edge devices are well-suited to execute the spatial attention map generation and initial decoupling gating, given their low dimensionality and compatibility with lightweight convolutional operators, which can be compiled for efficient inference on embedded neural accelerators under strict thermal budgets. Channel attention and phase-separation dynamics, which require more extensive memory and floating-point throughput, can be offloaded to fog nodes or centralized cloud clusters where dynamic resource scaling and renewable-powered data centers mitigate the carbon footprint. This split introduces a communication bottleneck, but the decoupled clothing embedding need not be transmitted if privacy policies mandate on-device suppression, simultaneously reducing bandwidth and improving sustainability by eliminating unnecessary data movement.

Sustainability metrics extend beyond operational energy to encompass the full lifecycle, including data collection, model training, and hardware refresh cycles. Training the dual-attention decoupling module on large-scale multi-domain re-identification datasets entails significant computational expense, encouraging the adoption of federated learning strategies where edge nodes collaboratively update a shared representation without pooling raw video data centrally. The phase-separation condensation module, being inherently incremental and local, enables online adaptation at the edge, curtailing the frequency of retraining and mitigating the carbon intensity associated with massive periodic cloud jobs. Lifecycle assessment methods borrowed from data center design can be applied to quantify the environmental break-even point between frequent localized condensation updates and occasional global model refresh, thus guiding operational policies. Moreover, hardware heterogeneity across deployment sites demands model compression techniques such as quantization and pruning tailored to the dual-attention architecture, preserving decoupling

fidelity while respecting device-specific power envelopes. The commitment to sustainability is thus woven into the system’s architecture, from modular compute partitioning to the algorithmic self-organization that reduces dependence on continuous heavy training.

## **6. Evaluation and Discussion**

Evaluating such a system requires a multidimensional framework that transcends conventional rank-1 accuracy. In addition to performance on standard and clothing-change benchmarks, one must assess cluster condensation integrity under temporal drift, measured by the stability of identity condensates when gallery populations evolve. A useful proxy is the mean condensation lifetime, reflecting how long an individual’s feature cluster remains distinct before merging or fragmenting. Robustness can be gauged through adversarial stress tests that inject synthetic appearance transformations mimicking urban seasonal changes or fashion cycles, while fairness audits employ disaggregated condensation metrics per demographic stratum to expose residual disparity. Interpretability evaluations can probe spatial attention maps for alignment with human-annotated identity-relevant keypoints, verifying that the system attends to body structure rather than superficial garment logos. Such assessments belong to a systems verification regimen that goes beyond task accuracy and addresses trustworthiness.

Comparisons with non-bio-inspired dual-attention systems reveal that while static decoupling approaches achieve promising results in controlled settings, they often exhibit brittle condensation in streaming deployments because they lack the self-correcting dynamics of the phase-separation module. The bio-inspired design thus demonstrates a pronounced advantage in longitudinal track-to-track association, where identity continuity must be maintained over days despite wardrobe changes. From an infrastructure perspective, the modular split-attention architecture showed easier integration into existing camera management software stacks, as its intermediate decoupled representations could be consumed by third-party analytics tools without exposing raw video, facilitating plug-and-play upgrades in smart city installations. Nevertheless, the complexity of tuning affinity thresholds for the condensation phase across heterogeneous environments remains a practical challenge; future work should explore meta-learning strategies that adaptively configure condensation hyperparameters based on environmental context detectors.

The broader implications of this work extend into domains such as wildlife re-identification and autonomous vehicle pedestrian tracking, where dynamic appearance changes and long-term identity association are equally critical. Translating the bio-inspired condensation principle to multispectral or event-based sensors would broaden robustness against lighting extremes, aligning with community efforts toward all-weather visual perception. Moreover, as regulatory frameworks evolve, the architectural transparency afforded by explicit decoupling can support certification for trustworthy AI under emerging standards, providing a reference implementation for how bio-inspiration can be deployed not merely for performance but for accountable autonomy.

## **7. Conclusion**

This paper presented a bio-inspired identity representation learning framework for robust visual re-identification, centered on dual-attention feature decoupling and cellular phase-separation dynamics. By treating identity encoding as an adaptive, self-organizing process rather than a static embedding, the system architecture addresses the persistent fragility of re-identification models under clothing changes and adversarial conditions. The dual-attention

module structurally separates identity-significant and appearance-related features, creating an auditable foundation for fairness, while the phase-separation-inspired condensation process stabilizes feature manifolds in a decentralized, context-sensitive manner. A systems-level analysis highlighted the trade-offs among modularity, privacy, latency, and energy efficiency, and positioned the design within a broader governance ecosystem that demands both technical robustness and ethical accountability. As identity-aware vision systems become woven into urban infrastructures, the integration of biological self-organization principles offers a pathway toward systems that are resilient, interpretable, and aligned with societal values.

## References

1. Luo, H., Gu, Y., Liao, X., Lai, S., & Jiang, W. (2019). Bag of tricks and a strong baseline for deep person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (pp. 0-0). IEEE.
2. Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., & Tian, Q. (2015). Scalable person re-identification: A benchmark. In Proceedings of the IEEE International Conference on Computer Vision (pp. 1116-1124). IEEE.
3. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
4. Yang, X., Zhang, T., Xu, C., Yan, S., Zhang, Y., & Tian, Q. (2021). Joint disentangling and adaptation for domain-generalizable person re-identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(7), 2490-2503.
5. Chung, C.-I., Yang, H., Jung, H., Lee, J., Kim, H., & Seong, R. H. (2024). Phase separation of YAP-MAML2 differentially regulates the transcriptome. *Proceedings of the National Academy of Sciences of the United States of America*, 121, e2310430121.
6. Woo, S., Park, J., Lee, J. Y., & Kweon, I. S. (2018). CBAM: Convolutional block attention module. In Proceedings of the European Conference on Computer Vision (ECCV) (pp. 3-19). Springer.
7. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., & Lu, H. (2019). Dual attention network for scene segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 3146-3154). IEEE.
8. Ding, Y., Wang, X., Yuan, H., Qu, M., & Jian, X. (2025). Decoupling feature-driven and multimodal fusion attention for clothing-changing person re-identification. *Artificial Intelligence Review*, 58(8), 241.
9. Chang, C., Fu, M., Chen, X., et al. (2025, November). Research on PDU-Net Lung Nodule Segmentation Algorithm Based on Path Aggregation and Dual Attention. In 2025 4th International Conference on Image Processing, Computer Vision and Machine Learning (ICICML) (pp. 1897-1900). IEEE.
10. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
11. Han, S., Mao, H., & Dally, W. J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In International Conference on Learning Representations (ICLR).

12. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645-3650). ACL.
13. Lyon, D. (2007). Surveillance studies: An overview. Polity.
14. Wang, G., Yuan, Y., Chen, X., Li, J., & Zhou, X. (2018). Learning discriminative features with multiple granularities for person re-identification. In Proceedings of the 26th ACM International Conference on Multimedia (pp. 274-282). ACM.
15. Li, D., Yang, Y., Song, Y. Z., & Hospedales, T. M. (2018). Learning to generalize: Meta-learning for domain generalization. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 32). AAAI Press.
16. Shin, Y., & Brangwynne, C. P. (2017). Liquid phase condensation in cell physiology and disease. *Science*, 357(6357), eaaf4382.
17. Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., & Stoica, I. (2010). Spark: Cluster computing with working sets. In Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing (HotCloud). USENIX Association.
18. Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In Proceedings of COMPSTAT'2010 (pp. 177-186). Physica-Verlag.
19. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE International Conference on Computer Vision (pp. 618-626). IEEE.
20. Chung, C.-I. et al. Phase separation of YAP-MAML2 differentially regulates the transcriptome. *Proc. Natl Acad. Sci. Usa.* 121, e2310430121 (2024).
21. Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In International Conference on Learning Representations (ICLR).