Knowledge-Guided Deep Learning Framework for Linking Medical Image Phenotypes and YAP-MAML2–Associated Transcriptional Dynamics in Cancer Progression Analysis

Pascal Ortiz
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
ortiz909@uab.edu

Felix C. Simpson
Department of Computer Science, University of North Texas, Denton, TX, USA.
simpson2001@unt.edu

Abstract

The integration of medical imaging and molecular profiling has been heralded as a cornerstone of precision oncology, yet the systems-level coupling between radiographic phenotypes and the transcriptional dynamics of oncogenic fusions remains poorly characterized. This paper presents a knowledge-guided deep learning framework designed to bridge the gap between image-based tumor phenotypes and the gene regulatory programs driven by the YAP-MAML2 fusion protein. The framework embeds curated pathway topologies, chromatin interaction maps, and condensate-mediated transcriptional models directly into a multimodal neural architecture, enforcing biological consistency while learning representations that connect morphological heterogeneity to downstream transcriptional trajectories. By structuring the architecture as a dual-stream system with shared latent alignments and constraint layers, the design hard-codes domain knowledge and offers an interpretable scaffold for probing the link between phenotypic imaging patterns and phase-separated transcriptional condensates. Drawing on lessons from large-scale biomedical AI deployments, the discussion extends beyond model design to encompass structural trade-offs, data governance, federated infrastructure, fairness, and long-term sustainability. The analysis underscores that the clinical translation of such knowledge-augmented systems depends as much on institutional interoperability, algorithmic equity, and regulatory alignment as on predictive accuracy. This work provides a systems-oriented blueprint for constructing, validating, and governing integrative deep learning systems that respect the complexity of both molecular biology and healthcare delivery.

Keywords

Knowledge-guided deep learning, medical image phenotyping, YAP-MAML2, transcriptional dynamics, cancer progression, multimodal fusion, system governance.

1. Introduction

The past decade has witnessed an unprecedented convergence of high-resolution medical imaging and high-throughput molecular profiling in oncology, enabling the co-analysis of tumor morphology, functional genomics, and transcriptional regulatory landscapes. Among the most compelling but underexplored connections is that between imaging phenotypes and

the spatiotemporal dynamics of oncogenic transcription factors, particularly those derived from fusion events that reprogram chromatin and gene expression programs. YAP-MAML2, a chimeric transcription factor recurrent in epithelioid hemangioendothelioma and other mesenchymal malignancies, exemplifies a class of fusions where phase separation behaviors create nuclear condensates that alter the three-dimensional enhancer landscape and drive aggressive transcriptional programs. Understanding how such molecular dynamics manifest in radiological and histopathological images could unlock noninvasive biomarkers of disease trajectory, yet the systems required to construct and deploy such linking models demand far more than simple correlative pipelines.

Current deep learning architectures have achieved remarkable success in segmenting and classifying medical images, but their capacity to link image features to molecular mechanisms remains limited when they operate in a purely data-driven fashion without enforcing biological plausibility. This limitation becomes especially acute when the target molecular phenotype is not a static gene expression signature but a dynamic transcriptional trajectory governed by phase-separated assemblies of fusion proteins, nucleosome remodeling, and enhancer switching. A knowledge-guided learning paradigm, in which curated biological graphs, pathway constraints, and mechanistic models are woven into the architecture and loss functions, offers a principled route to building models that are simultaneously predictive and consistent with established molecular biology. Such a paradigm shifts the role of prior knowledge from post hoc interpretation to an active structural component of model design, thereby addressing generalization brittleness and improving robustness across heterogeneous patient populations [1, 2].

This paper proposes a system-level deep learning framework that links medical image phenotypes with YAP-MAML2–associated transcriptional dynamics by integrating multimodal inputs, hierarchical attention mechanisms, and knowledge-derived priors. The architecture is conceptualized not merely as a computational method but as a sociotechnical infrastructure that must contend with data heterogeneity, privacy constraints, deployment scalability, and ethical governance. Drawing from earlier efforts in radiogenomics, knowledge-guided graph neural networks, and federated learning, the framework formalizes the connection between imaging feature spaces and gene regulatory networks through a shared latent embedding that is regularized by known pathway relationships and condensate biology. In the following sections, we unpack the components of this framework, analyze the structural and operational trade-offs inherent in its design, and situate the system within the broader landscape of translational AI governance.

2. Background and Related Work

Medical image phenotyping has evolved from hand-crafted radiomics to learned representations extracted by deep convolutional and transformer-based models. Foundational segmentation architectures such as U-Net demonstrated that end-to-end trainable networks could capture hierarchical image features suitable for delineating tumor boundaries and characterizing tissue texture [4]. Subsequent advances in radiomics formalized the extraction of engineered image features as a bridge to molecular correlates, although such features often failed to capture the complex, non-linear interactions present in biological systems [5]. More recent self-supervised and multimodal models have begun to jointly embed images and molecular profiles, yet they typically treat the molecular output as a static snapshot rather than a dynamic process shaped by transcriptional condensates and chromatin remodeling [20].

Parallel to these imaging advances, the study of oncogenic fusion proteins has revealed that many drivers, including YAP-MAML2, function through phase separation mechanisms that create condensates at super-enhancer regions, rewiring transcriptional programs in ways that cannot be captured by simple differential expression analyses [3, 7]. Chung and colleagues demonstrated that YAP-MAML2 forms concentration-dependent liquid-like droplets that differentially regulate target genes depending on the structural integrity of the fusion's interaction domains [3]. These findings align with broader models of transcriptional control in which phase-separated condensates concentrate transcription factors and coactivators to modulate gene expression bursting [8]. Yet the link between such molecular dynamics and the macroscopic tumor architecture visible in computed tomography or magnetic resonance imaging remains largely unexplored.

Attempts to bridge imaging and molecular data have often employed multimodal fusion strategies that concatenate or co-attend features from different modalities [10]. While these approaches can improve predictive performance, they rarely enforce consistency with gene regulatory network topologies. Knowledge graph-based methods have been introduced to incorporate structured biomedical knowledge, such as protein–protein interactions and pathway hierarchies, into deep learning pipelines [9]. Similarly, graph convolutional networks have been applied to model polypharmacy side effects by leveraging molecular interaction networks, demonstrating that explicit knowledge representations enhance generalization [19]. However, applying such strategies to the imaging-transcriptional dynamics axis requires additional infrastructure capable of aligning heterogeneous data types, handling missing modalities, and modeling temporal dependencies that emerge from transcriptional feedback loops.

Previous work in radiogenomics has highlighted the prognostic value of integrating imaging features with gene expression signatures, but most efforts have focused on static pathway enrichment scores rather than dynamic, mechanism-informed trajectories [6, 23]. The framework proposed here extends these ideas by embedding a curated model of YAP-MAML2 condensate formation, TEAD-dependent and independent regulatory circuits, and knowledge of chromatin looping into the training process itself, thereby constraining the learned latent space to reflect known biology while allowing the image-derived features to modulate the state of the transcriptional system.

3. System Architecture and Knowledge Integration Paradigm

The proposed framework is built around a dual-stream architecture that processes imaging data and transcriptional data through separate encoder branches before fusing them in a shared bottleneck that is regularized by domain knowledge. The image stream accepts multiscale volumetric inputs, such as contrast-enhanced CT or whole-slide histopathology images, and applies hierarchical feature extraction using a combination of residual convolutional blocks and vision transformer layers. The transcriptional stream ingests time-course RNA expression profiles or pseudotime-ordered single-cell transcriptomic data and encodes them through a recurrent neural network constrained by prior gene regulatory network structures. The two streams are aligned via a cross-modal attention mechanism that learns a multimodal latent representation of tumor state, from which transcriptional trajectory predictions are generated.

Crucially, the framework injects prior biological knowledge at multiple levels. At the input stage, gene expression features are grouped and hierarchically organized according to pathway databases and protein complex membership, forcing the network to respect

functional modularity. In the latent fusion space, a knowledge-based regularization term penalizes deviations from known causal relationships encoded in a curated knowledge graph that connects imaging-derived features, such as necrosis patterns or vascular heterogeneity, to transcription factor activities via intermediate signaling nodes. Additionally, a soft constraint layer encodes the biophysical principles of YAP-MAML2 phase separation, including the cooperative binding dynamics and concentration thresholds that govern condensate formation, as a differentiable module that influences the transition function of the transcriptional decoder. This design shifts the learning problem from a purely empirical risk minimization to a constrained optimization problem where biological feasibility acts as a structural prior, increasing sample efficiency and reducing overfitting to spurious correlations that plague purely data-driven radiogenomic models [21, 2].

The architectural trade-offs inherent in such a design are substantial. Integrating rigid knowledge constraints risks reducing model capacity to discover novel patterns not yet captured in curated databases, a tension familiar from symbolic-subsymbolic hybrid systems. The framework mitigates this by adopting a soft constraint formulation with tunable penalty coefficients that can be annealed over training, allowing the network to initially prioritize reconstruction accuracy before gradually tightening biological consistency. On the deployment side, the dual-stream architecture with cross-modal attention introduces computational overhead that challenges real-time inference in clinical settings, motivating the use of knowledge distillation to compress the trained model into a lightweight student network suitable for edge deployment [17]. The residual connections and attention layers must also be carefully regularized to prevent mode collapse when one modality is missing, a common scenario in clinical workflows.

4. Data Integration and Multimodal Phenotype Representation

Building a system that associates radiological phenotypes with transcriptional dynamics requires addressing the profound heterogeneity of data sources, acquisition protocols, and scales. Imaging data may include CT scans reconstructed with variable kernel parameters, histopathology slides stained under different conditions, and time-series imaging from longitudinal studies. Transcriptional data, in contrast, may arise from bulk tumor RNA-seq at a single time point, spatial transcriptomics, or single-cell pseudotime trajectories that approximate dynamic progression. Aligning these disparate data modalities into a common representation space demands a careful orchestration of preprocessing, representation learning, and domain adaptation mechanisms that go beyond simple concatenation.

The framework's image stream benefits from transfer learning from pre-trained self-supervised models on large-scale chest CT repositories, which capture domain-agnostic textural and structural features that are then fine-tuned on smaller annotated cohorts [20]. For lung lesion segmentation, state-of-the-art approaches that use path aggregation and dual attention mechanisms have demonstrated robust boundary delineation even in heterogeneous nodule populations, providing high-fidelity morphological masks that feed downstream radiomic and deep feature extraction [16]. The resulting feature vectors encode not only size and shape but also internal texture heterogeneity, peritumoral halo characteristics, and enhancement kinetics when dynamic imaging is available. These features are passed through a domain adaptor network that learns scanner-invariant representations using adversarial training, thereby reducing the risk that site-specific artifacts will dominate the learned latent mapping to transcriptional states [18].

On the transcriptional side, the framework models expression data as a function of a latent transcriptional state vector that evolves according to regulatory constraints. Single-cell RNA velocity and pseudotime inference provide a coarse temporal ordering that permits the construction of trajectories along which the activity of YAP-MAML2 condensates can be traced. By pooling cells along inferred trajectories and projecting them into a shared low-dimensional space with the imaging features using multi-omics factor analysis–inspired alignment techniques, the system establishes a joint latent manifold where proximity encodes both morphological and molecular similarity [23]. This fusion step is critical for enabling the cross-modal attention mechanism to selectively weight the influence of specific image regions on particular gene modules, a process that would be uninterpretable without the alignment.

The multimodal representation raises governance challenges related to data provenance, informed consent, and institutional interoperability. Because image and molecular data are frequently housed in separate clinical and research databases, linkage requires robust hashing and de-identification protocols that preserve privacy while allowing cross-modality mapping. Federated learning architectures that keep data localized and share only encrypted model updates during training represent a structurally aligned solution, but they introduce additional coordination overheads and heterogeneity in local computational resources that must be considered in any deployment roadmap [12].

5. Transcriptional Dynamics Modeling and Knowledge-Guided Learning

The core scientific objective of the framework is to predict how YAP-MAML2–driven transcriptional programs evolve as a function of tumor morphological changes observable in medical images. YAP-MAML2 forms nuclear condensates through phase separation, and the resultant condensates congregate at enhancer regions marked by TEAD and other cofactors, altering the bursting frequency of target genes and thereby reshaping the cellular transcriptome over time [3, 7]. Modeling these dynamics requires a transition model that respects the hysteretic behavior of condensate assembly and disassembly, the competitive binding among coactivators, and the nonlinearities introduced by chromatin looping.

To encapsulate this biology, the framework incorporates a constrained recurrent unit whose state transition parameters are derived from a biophysical phase separation model. Specifically, the latent vector representing transcriptional activity is updated in discrete time steps according to a function that is regularized to reflect the concentration-dependent formation and dissolution of condensates as described by the Flory-Huggins theory of polymer phase separation, although the exact equations are not exposed at the level of mathematical notation within this discussion. The knowledge-guided component comes from the observation that the gene targets of YAP-MAML2 are not random but are organized into co-regulated modules identified through systematic gene network inference and chromatin immunoprecipitation experiments [6, 9]. By hard-coding the membership of genes into these modules and using attention masks derived from Hi-C and ChIA-PET data, the model is forced to predict transcriptional changes only in gene neighborhoods that physically interact with condensate-occupied enhancers.

The training process leverages both supervised and self-supervised objectives. The supervised loss compares predicted expression profiles from the transcriptional decoder with ground truth RNA-seq data at matched time points or pseudotemporal positions. A self-supervised contrastive loss encourages the image-derived latent state of tumors with similar condensate gene module activity to occupy nearby positions in the joint embedding, even when the imaging appearance diverges due to unrelated stromal reactions. This multiple-objective

training schedule, combined with adaptive knowledge-constraint weighting, allows the system to learn fine-grained mappings that associate, for instance, a heterogeneous rim enhancement pattern on CT with a transition from a transcriptionally quiescent condensate state to a highly active super-enhancer engagement phase. Such associations are not predetermined but emerge from the interplay of data and prior knowledge, and they can subsequently be probed using interpretability tools such as Grad-CAM for image regions and SHAP values for gene contributions [11, 22].

Interpreting the learned links between imaging phenotypes and transcriptional states is essential for building clinical trust and generating testable biological hypotheses. The framework's explicit knowledge layers act as a built-in explanation module, since each node in the knowledge graph can be traced back to published pathway annotations. When deployed prospectively, the system can produce not only a predicted transcriptional trajectory but also a relevance map highlighting the most influential imaging features and the corresponding intermediate molecular nodes, thus enabling clinicians and biologists to evaluate the plausibility of the prediction against known oncogenic mechanisms.

6. Infrastructure, Scalability, and Deployment Trade-offs

Translating a knowledge-guided multimodal deep learning framework from a research prototype to a clinical-grade system entails confronting a set of infrastructure and scalability challenges that are as formidable as the modeling problems themselves. Training the full dual-stream architecture with dynamic knowledge constraints demands significant GPU memory and data throughput, especially when handling volumetric imaging and large gene expression matrices simultaneously. Distributed training across multiple institutions, which is often necessary to amass sufficiently diverse and rare fusion-positive cases, introduces network latency, heterogeneous hardware, and the need for fault-tolerant synchronization protocols.

Federated learning has emerged as a principled solution that allows collaborative model training without exposing patient-level data. In a federation of cancer centers, each site can train a local copy of the framework on its own imaging and molecular data, exchanging only model gradients or parameter updates through a secure aggregation server. The architecture's modular design, with separate image and transcriptional encoders, facilitates partial federation where sites that lack in-house molecular profiling capabilities can still contribute by training the image encoder branch using a shared frozen decoder provided by a coordinating node. This partial contribution model aligns with the reality of uneven molecular testing infrastructure across healthcare systems and promotes broader participation [12]. Nonetheless, federated training introduces governance complexities related to liability when models trained on multi-jurisdictional data are subsequently deployed in local settings, requiring careful auditing trails and version control.

The inference stage demands that the system operate within the time constraints of clinical radiology and oncology workflows, which typically require results within seconds for integrated reporting. A full dual-stream forward pass with cross-modal attention on high-resolution volumes is computationally intensive, but the soft constraint module and attention layers can be aggressively pruned or distilled into a more compact model without losing the critical knowledge-guided structure. Knowledge distillation, wherein a smaller student network is trained to match the predictions and internal knowledge representations of the larger teacher model, has proven effective in reducing memory and latency while preserving the inductive biases imparted by the knowledge constraints [17]. When combined with

hardware accelerators optimized for medical imaging inference, such a distilled model can be deployed at the point of care, enabling real-time interactive analysis.

Sustainability and energy efficiency are increasingly recognized as core design goals for large-scale AI systems in healthcare. The computational cost of training extensive multimodal models, particularly those with recurrent dynamics and adversarial domain adaptation, must be weighed against their marginal clinical benefit. The framework's integration of explicit biological priors serves an efficiency purpose by reducing the effective search space of network parameters, potentially decreasing the number of training epochs required and thereby lowering the carbon footprint relative to unconstrained foundation models [25]. Choosing a minimal necessary complexity and incorporating periodic model pruning are operational strategies that infrastructure architects can adopt to align environmental responsibility with clinical effectiveness.

7. Governance, Fairness, and Ethical Implications

The deployment of a system that connects medical images to molecular dynamics carries profound governance and ethical responsibilities that extend far beyond technical validation. Radiogenomic models can inadvertently learn spurious associations that reflect biases in training data, such as varying scanner technology, institutional referral patterns, and demographic underrepresentation. If not explicitly mitigated, these biases can lead to systematically different predictive accuracy across population subgroups, raising fairness concerns that have both regulatory and moral dimensions. Studies have shown that models trained predominantly on data from one sex or ethnic group can underperform when applied to others, a risk amplified in rare fusion-driven cancers where cohort sizes are small and demographic skew is pronounced [13].

The framework addresses fairness through several complementary mechanisms. The knowledge-guided constraints inherently regularize the model toward biologically conserved relationships that transcend population variations, damping the influence of spurious correlations that might arise from local imaging artifacts. In addition, the domain adaptation layers described earlier reduce scanner-specific biases, while stratified federated aggregation can ensure that no single population dominates the global model update. Auditing tools that report performance disaggregated by sex, age, and self-identified racial categories are integral to the governance model, and continuous post-market monitoring should be implemented following regulatory guidance on adaptive AI systems in medicine [14].

Regulatory pathways for such integrated systems remain evolving. As a software-based medical device that combines image analysis with molecular prediction, the framework would likely be classified by agencies such as the U.S. Food and Drug Administration under a risk-based framework that scrutinizes the intended use, the strength of the clinical evidence linking predicted transcriptional states to actionable outcomes, and the transparency of the underlying model. The inclusion of interpretable knowledge layers facilitates the production of detailed documentation for regulatory review, because every major model decision can be contextualized within established biological knowledge graphs. This stands in contrast to black-box foundation models, for which explainability remains an open challenge. However, the dynamic nature of knowledge bases—where new biological discoveries can alter previously accepted pathways—creates a versioning challenge that demands a rigorous change control and revalidation process analogous to that used for clinical decision support algorithms.

Data governance further intersects with privacy when imaging and genomic data are linked. Even when identifiers are stripped, the combination of a medical image and a transcriptional profile can create a unique re-identification risk, as has been demonstrated for genomic data alone [24]. Institutional data use agreements must specify the purpose and duration of multimodal linking, and patients must be informed about the potential inferences that could be drawn from their combined data. The architecture's federated design reduces the concentration of linkable data in a single repository, but robust cryptographic protocols, including differential privacy guarantees, are advisable for the aggregation step. Policymakers and institutional review boards need to develop clear rubrics for evaluating the trade-off between the societal benefit of understanding rare cancer fusions and the individual risk of re-identification.

8. Conclusion

The proposed knowledge-guided deep learning framework represents a systems-level approach to connecting medical image phenotypes with YAP-MAML2–associated transcriptional dynamics, moving beyond correlative radiogenomics toward a mechanistic integration of morphological data and phase-separated transcriptional regulatory programs. By embedding curated biological knowledge directly into a dual-stream multimodal architecture, the framework addresses long-standing challenges of interpretability, sample efficiency, and biological consistency that have hindered the clinical translation of integrative cancer models. The analysis has made clear that the value of such a system lies not only in its predictive capability but in its capacity to serve as a testable hypothesis generator that aligns imaging readouts with the molecular grammar of oncogenic fusion proteins.

Equally important is the recognition that the successful deployment of such a system is contingent on navigating a complex landscape of infrastructure scalability, federated coordination, algorithmic fairness, and regulatory compliance. The trade-offs between model complexity and inference speed, between rigid biological constraints and openness to novel discovery, and between centralized governance and federated autonomy must be deliberately balanced. The framework described here offers a reference design that can be iterated upon as the underlying biological knowledge evolves and as healthcare systems mature in their digital transformation. Through careful attention to governance structures, sustainability imperatives, and ethical data practices, it is possible to build not merely accurate models but trustworthy and equitable systems that advance precision oncology for rare fusion-driven malignancies.

References

1. Acosta, J. N., Falcone, G. J., Rajpurkar, P., & Topol, E. J. (2022). Multimodal biomedical AI. Nature Medicine, 28, 1773–1784.

2. Al-Shaer, L., Buddenkotte, T., & Schönlieb, C.-B. (2022). Knowledge-guided deep learning of breast cancer histopathology images using gene expression and pathway annotations. Scientific Reports, 12, 17438.

3. Chung, C. I., Yang, J., Yang, X., Liu, H., Ma, Z., Szulzewsky, F., ... & Shu, X. (2024). Phase separation of YAP-MAML2 differentially regulates the transcriptome. Proceedings of the National Academy of Sciences, 121(7), e2310430121.

4. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015 (pp. 234–241). Springer.

5. Lambin, P., Leijenaar, R. T. H., Deist, T. M., Peerlings, J., de Jong, E. E. C., van Timmeren, J. E., ... & Aerts, H. J. W. L. (2017). Radiomics: the bridge between medical imaging and personalized medicine. Nature Reviews Clinical Oncology, 14(12), 749–762.

6. Marbach, D., Costello, J. C., Küffner, R., Vega, N. M., Prill, R. J., Camacho, D. M., ... & Kellis, M. (2012). Wisdom of crowds for robust gene network inference. Nature Methods, 9(8), 796–804.

7. Papoudou-Bai, A., Barbouti, A., Galani, V., Stefanaki, K., & Kanavaros, P. (2020). YAP1-MAML2 fusion in cancer: molecular pathogenesis and therapeutic implications. International Journal of Molecular Sciences, 21(13), 4700.

8. Hnisz, D., Shrinivas, K., Young, R. A., Chakraborty, A. K., & Sharp, P. A. (2017). A phase separation model for transcriptional control. Cell, 169(1), 13–23.

9. Himmelstein, D. S., Lizee, A., Hessler, C., Brueggeman, L., Chen, S. L., Hadley, D., Green, A., Khankhanian, P., & Baranzini, S. E. (2017). Systematic integration of biomedical knowledge prioritizes drugs for repurposing. eLife, 6, e26726.

10. Huang, S. C., Pareek, A., Seyyedi, S., Banerjee, I., & Lungren, M. P. (2020). Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. NPJ Digital Medicine, 3, 136.

11. 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).

12. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Maier-Hein, K. (2020). The future of digital health with federated learning. NPJ Digital Medicine, 3, 119.

13. Cirillo, D., Catuara-Solarz, S., Morey, C., Guney, E., Subirats, L., Mellino, S., ... & Vellido, A. (2020). Sex and gender differences and biases in artificial intelligence for biomedicine and healthcare. NPJ Digital Medicine, 3, 81.

14. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of artificial intelligence technologies in medicine. Nature Medicine, 25, 30–36.

15. Marusyk, A., & Polyak, K. (2010). Tumor heterogeneity: causes and consequences. Biochimica et Biophysica Acta (BBA)-Reviews on Cancer, 1805(1), 105–117.

16. Chang, C., Fu, M., Chen, X., Feng, S., Zhang, M., Zhou, X., ... & Liu, Z. (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.

17. Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531.

18. Tzeng, E., Hoffman, J., Saenko, K., & Darrell, T. (2017). Adversarial discriminative domain adaptation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 7167–7176).

19. Zitnik, M., Agrawal, M., & Leskovec, J. (2018). Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics, 34(13), i457–i466.

20. Azizi, S., Mustafa, B., Ryan, F., Beaver, Z., Freyberg, J., Deaton, J., ... & Norouzi, M. (2021). Big Self-Supervised Models Advance Medical Image Classification. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 3458–3468).

21. Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. Proceedings of the IEEE, 109(5), 612–634.

22. Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., ... & Lee, S. I. (2020). From local explanations to global understanding with explainable AI for trees. Nature Machine Intelligence, 2, 56–67.

23. Argelaguet, R., Velten, B., Arnol, D., Dietrich, S., Zenz, T., Marioni, J. C., ... & Stegle, O. (2018). Multi-omics factor analysis—a framework for unsupervised integration of multi-omics data sets. Molecular Systems Biology, 14, e8124.

24. Gymrek, M., McGuire, A. L., Golan, D., Halperin, E., & Erlich, Y. (2013). Identifying personal genomes by surname inference. Science, 339, 321–324.

25. 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).