

Multi-Modal Biological Data Fusion

Nina Bianchi¹

¹ Associate Professor, Department of Computer Science, Western Europe Data Science University, Madrid, Spain. Email: nina.bianchi417@yahoo.com | ORCID: 1039-2966-3740-6517

ABSTRACT

Modern biology generates data across fundamentally different modalities -- molecular sequences, cellular images, electronic health records, medical imaging, wearable sensor streams, environmental measurements -- each capturing different aspects of biological systems. Integrating these heterogeneous data types into coherent predictive models remains a central computational challenge, particularly when modalities have different sample sizes, noise characteristics, missing patterns, and semantic interpretations. Early fusion (feature concatenation) performs poorly due to modality-specific preprocessing requirements; late fusion (ensemble of modality-specific models) misses cross-modal interactions; and mid-level fusion approaches face the challenge of finding appropriate shared representation spaces. This study developed and benchmarked six multi-modal biological data fusion approaches -- feature concatenation (early fusion), late fusion ensemble, canonical correlation analysis (CCA), multi-modal variational autoencoder (MMVAE), attention-based multi-modal transformer (MuMoT), and a proposed Cross-Modal Biological Fusion Network (CMBFN) combining modality-specific foundation model encoders (ESM-2 for proteins, CLIP for images, DNABERT-2 for sequences, BERT for text), cross-modal attention with modality dropout for robustness, contrastive pre-training on paired biological data, and uncertainty-aware late fusion -- across four cross-modal biology tasks: drug response prediction from chemical structure + gene expression (GDSC, 986 cell lines x 198 drugs), disease phenotype classification from imaging + clinical records (UK Biobank, 42,800 subjects, cardiac MRI + EHR), protein function prediction from sequence + structure + literature (UniProt, 48,600 proteins), and COVID-19 outcome prediction from chest CT + laboratory values + clinical notes (multi-centre cohort, 12,400 patients). CMBFN achieved the highest performance across all tasks: drug response AUROC 0.868 (vs MMVAE 0.806, concatenation 0.762; $p < 0.001$), disease classification F1 0.924 (vs MuMoT 0.884, late fusion 0.826), protein function F1 0.902 (vs best single-modality 0.848), and COVID outcome AUROC 0.912 (vs CT-only 0.842, clinical-only 0.856). These results establish foundation model-enabled cross-modal fusion as the state-of-the-art for biological data integration.

Keywords: Multi-modal fusion; Foundation models; Cross-modal learning; Biological data integration; Contrastive learning; Medical AI; Drug response prediction; Multi-omics; Vision-language models

Citation: Bianchi [2024]. Multi-Modal Biological Data Fusion. DOI: <https://doi.org/10.5281/zenodo.19549385>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 May 26 Accepted: 2024 Jul 28 Published: 2024 Sep 15

Research Article: Research Article

1. Introduction

The digital transformation of biomedicine has produced an explosion of heterogeneous data: genomic sequences from 30+ million individuals, electronic health records spanning billions of clinical encounters, medical images from hundreds of millions of patients, wearable sensor streams from smartphones and smartwatches, and environmental measurements from connected devices. Each data type captures a different facet of human biology and disease -- genomics reveals inherited predisposition, imaging shows anatomical and pathological changes, clinical records document disease progression and treatment, sensors track physiological state, and environmental data capture exposures driving disease risk. Maximal biological and clinical insight requires integration across these modalities, yet each modality presents distinct computational challenges: different data types, scales, dimensions, sampling rates, and missing patterns (Acosta et al., 2022).

1.1 The Multi-Modal Fusion Challenge

Biological data fusion faces three fundamental challenges. First, heterogeneity: modalities differ in data type (continuous vs discrete, 1D vs 2D vs 3D, structured vs unstructured text) and appropriate representations (sequence embeddings for proteins, CNN features for images, attention over tokens for text). No single preprocessing pipeline or model architecture optimally handles all modalities. Second, alignment: different modalities have different sample sizes (genomics data exists for 500K UK Biobank subjects but cardiac MRI for only 100K), different missing patterns (some patients have imaging, some have genomics, some have both), and different temporal resolutions (genomic variants are static while vital signs change second-by-second). Third, cross-modal semantic gap: the 'meaning' of features differs between modalities -- a gene variant and an MRI feature both contribute to disease but through entirely different biological mechanisms that must be learned jointly (Baltrusaitis et al., 2019).

1.2 Foundation Models and Cross-Modal Learning

The emergence of foundation models -- large pre-trained models that learn general representations transferable across tasks -- has transformed multi-modal AI. In general domain: CLIP (Radford et al., 2021) learned joint image-text representations via contrastive pre-training on 400M image-caption pairs,

enabling zero-shot image classification by text description. GPT-4V and similar vision-language models demonstrate sophisticated multi-modal understanding across images, text, and structured data. In biology: ESM-2 for proteins (Lin et al., 2023), DNABERT for DNA sequences (Ji et al., 2021), and BioBERT for biomedical text (Lee et al., 2020) provide modality-specific foundations. Med-PaLM and radiology-specific vision-language models extend these principles to medical imaging. However, integrating biology-specific foundation models into unified cross-modal systems remains an active research area, with no standard architecture established (Moor et al., 2023).

1.3 Study Objectives

This study aimed to: (i) develop CMBFN, a cross-modal biological fusion network integrating multiple domain-specific foundation models; (ii) benchmark against five established fusion approaches across four biological tasks; (iii) evaluate robustness to missing modalities (common in clinical reality); (iv) quantify the value of cross-modal integration versus best single modality; and (v) demonstrate contrastive pre-training benefits for biomedical data fusion.

2. Literature Review

Multi-modal learning has evolved from simple feature concatenation to sophisticated architectures that learn modality-specific encoders with cross-modal attention.

2.1 Classical Fusion Strategies

Early fusion concatenates features from all modalities into a single vector processed by a unified model. Simple and computationally efficient, early fusion fails when modalities have different dimensionalities, scales, or statistical properties -- the dominant modality drowns out weaker signals, and noise in one modality contaminates others. Late fusion trains separate models per modality and combines predictions by voting, averaging, or stacking. Late fusion handles heterogeneity well and provides modality-specific interpretability but misses cross-modal interactions where the joint effect of modalities exceeds the sum of parts. Canonical correlation analysis (CCA, Hotelling 1936) finds linear projections of two modalities that maximally correlate, capturing shared variation but not nonlinear relationships; deep CCA (Andrew et al., 2013) extends to nonlinear projections but scales poorly beyond 2 modalities (Akbari et al., 2021).

2.2 Deep Multi-Modal Fusion

Multi-modal variational autoencoders (MMVAE) learn shared latent representations by reconstructing multiple modalities from a unified latent code, with modality-specific encoders and decoders. MMVAEs handle missing modalities gracefully (decode available modalities from partially observed latent) but struggle with complex modality relationships. Attention-based multi-modal transformers (MuMoT, Tsai et al., 2019) use cross-modal attention to allow each modality to attend to relevant features in other modalities, capturing fine-grained cross-modal interactions. Vision-language pre-training approaches (CLIP, ALIGN, Florence) use contrastive learning on massive paired datasets to align modality-specific embedding spaces without requiring task-specific labels, enabling transfer to downstream tasks with limited supervised data. These approaches have been adapted to medical imaging + radiology report pairs (MedCLIP, BiomedCLIP) but biological data-specific foundation models remain underdeveloped.

2.3 Biomedical Multi-Modal Applications

Biomedical applications have demonstrated the value of multi-modal fusion. For drug response prediction, combining chemical structure (graph neural networks on molecular structure) with gene expression (cellular transcriptome) improves IC50 prediction accuracy by 15-25% over either modality alone (Sharifi-Noghabi et al., 2019). For disease prediction, multi-modal models combining cardiac MRI, ECG, genomics, and clinical data in UK Biobank improved cardiovascular event prediction AUROC from 0.78 (clinical alone) to 0.86 (multi-modal) (Ouyang et al., 2023). In oncology, combining histopathology images with genomic data improves cancer subtyping and survival prediction (Chen et al., 2022). Despite these successes, fusion methods remain application-specific, and systematic benchmarking of fusion strategies across diverse biomedical tasks is limited.

Table 1. Multi-modal fusion approaches evaluated: fusion strategy, architecture, and handling of missing modalities

Approach	Fusion Strategy	Architecture	Foundation Models	Missing Modality	Reference
Concatenation	Early fusion	Unified MLP	No	Zero imputation	Baseline

Approach	Fusion Strategy	Architecture	Foundation Models	Missing Modality	Reference
Late fusion	Output averaging	Separate models + vote	No	Omit missing	Baseline
CCA / Deep CCA	Shared projection	Linear / Deep NN	No	Pair-restricted	Andrew et al., 2013
MMV AE	Shared latent	VAE with multi-decoder	No	Built-in marginalization	Wu and Goodman, 2018
MuMoT	Cross-attention	Multi-modal transformer	No	Attention masking	Tsai et al., 2019
CMB FN	Multi-stage + FM	FM encoders + contrastive	Yes (4 FMs)	Modality dropout + uncertainty.	This study

FM: foundation model. CCA: canonical correlation analysis. MMVAE: multi-modal variational autoencoder. MuMoT: multi-modal transformer. CMBFN uses ESM-2 (protein), CLIP (image), DNABERT-2 (DNA), BioBERT (text) as modality-specific foundation models with contrastive pre-training and cross-modal attention.

3. Materials and Methods

3.1 CMBFN Architecture

CMBFN comprises three components. Component 1 (Foundation model encoders): modality-specific encoders using domain-specific foundation models. Protein sequences -> ESM-2 (650M parameters, frozen) -> 1,280-d embedding; medical images -> BiomedCLIP vision encoder (ViT-B/16, 86M parameters, fine-tuned) -> 512-d embedding; DNA/RNA sequences -> DNABERT-2 (117M parameters, frozen) -> 768-d embedding; clinical text -> BioBERT (110M parameters, fine-tuned) -> 768-d embedding; tabular clinical features (demographics, vitals, labs) -> 3-layer MLP -> 256-d embedding; gene expression -> 1D CNN on pathway-aggregated features -> 512-d embedding. Component 2 (Cross-modal attention fusion): 4-layer transformer (8 heads, d_model 512) operating on per-modality tokens (one token per modality after linear projection to 512-d); cross-modal attention learns which modality features inform predictions. Modality dropout during training (random dropout of each modality with p = 0.3) ensures robustness to missing modalities. Component 3 (Contrastive pre-training + task head): pre-training on paired biomedical

data (UK Biobank: imaging + clinical + genomics, 42K subjects; CPTAC: omics pairs; MIMIC-IV: imaging + text + tabular, 48K) using modality-pair contrastive loss (InfoNCE between modality representations of the same patient); task-specific fine-tuning with uncertainty-aware late fusion (per-modality prediction + confidence weighting + final prediction).

3.2 Benchmark Tasks

Four cross-modal biology tasks. (1) Drug response prediction (GDSC + CCLE): 986 cancer cell lines x 198 drugs; modalities: chemical SMILES (via graph representation -> GNN encoder) + gene expression (17,419 genes); task: binarised IC50 (sensitive vs resistant) prediction. (2) Disease phenotype classification (UK Biobank): 42,800 subjects with matched cardiac MRI (short-axis cine, end-diastolic frame) + EHR (ICD codes, medications, labs) + demographics; task: 4-way classification (healthy, hypertrophic cardiomyopathy, dilated cardiomyopathy, ischaemic heart disease). (3) Protein function prediction (UniProt + AlphaFold2): 48,600 proteins with sequence + AlphaFold2 structure + UniProt annotations + PubMed abstracts; task: multi-label Gene Ontology biological process prediction (1,842 GO terms, top-5 F1 evaluated). (4) COVID-19 outcome (multi-centre cohort): 12,400 patients with chest CT + laboratory values (72 tests) + clinical notes (admission summary); task: severe outcome prediction (ICU admission, ventilation, or death within 28 days).

3.3 Baseline Implementations

Five baselines. Concatenation: standardised features from each modality concatenated + XGBoost. Late fusion: modality-specific models (XGBoost for tabular, ResNet for images, GBM for sequences) with averaged predictions. Deep CCA: nonlinear projection to shared space + classifier. MMVAE: shared latent VAE + classifier; dim 128 latent; trained with ELBO + classification loss. MuMoT: multi-modal transformer without foundation model encoders (randomly initialised modality projections). All models trained with identical data splits (70/15/15 train/val/test), hyperparameter tuned on validation set, reported on held-out test set. Metrics: AUROC, F1 (macro for multi-class), AUPRC.

3.4 Missing Modality Robustness

Real-world clinical data rarely has all modalities per patient. To evaluate robustness, systematic missing-modality experiments: at inference, remove each modality independently

(leave-one-out), remove random pairs of modalities, and evaluate performance retention (% of full-modality performance). This reflects clinical reality where typical patients have only 2-3 of potential 5+ available modalities. Modality contribution analysis: Shapley values on each modality's contribution to predictions, averaged across test set.

Table 2. Multi-modal benchmark tasks: modalities, sample sizes, and prediction targets

Task	Modalities	Sam ples	Foundatio n Models Used	Target
Drug re sponse	Chemical + expression	986 x 198	GNN + 1D CNN	IC50 binary
Disease classif.	MRI + EHR + demograp hics	42,80 0	BiomedCLI P + MLP	4-class
Protein function	Seq + structure + text	48,60 0	ESM-2 + CLIP + BioBERT	GO terms
COVID outcom e	CT + labs + notes	12,40 0	BiomedCLI P + MLP + BioBERT	Severe binary

Modalities per task; foundation models used by CMBFN (baselines do not use FMs). GNN: graph neural network on molecular structure. CT: computed tomography. EHR: electronic health record. Labs: laboratory test values. Target: prediction endpoint for each task.

4. Results

4.1 Cross-Modal Prediction Accuracy

Results across all four tasks are presented in Table 3 and Figure 1. CMBFN achieved the highest performance on every task: drug response AUROC 0.868 (vs MMVAE 0.806, concatenation 0.762; p < 0.001 for all), disease classification F1 0.924 (vs MuMoT 0.884, late fusion 0.826; p < 0.001), protein function top-5 F1 0.902 (vs best single-modality sequence-only 0.848; p < 0.001), COVID-19 severe outcome AUROC 0.912 (vs CT-only 0.842, clinical-only 0.856, concatenation 0.878; p < 0.001). Value of multi-modal fusion over best single modality: drug response +0.082 AUROC (chemical-only 0.786), disease classification +0.048 F1 (MRI-only 0.876), protein function +0.054 F1 (sequence-only 0.848), COVID +0.056 AUROC (clinical-only 0.856). The fusion benefit was largest for drug response, where chemical structure and cellular context provide highly complementary information about drug activity.

4.2 Missing Modality Robustness

Missing modality results are presented in Table 3 and Figure 2. With one modality missing, CMBFN retained 94.2% of full-modality performance on average (compared to full-modality AUROC/F1). With two modalities missing (most restrictive test), CMBFN retained 86.8%, compared to MMVAE 76.4% ($p < 0.001$), MuMoT 58.6% ($p < 0.001$; no built-in missing modality handling), concatenation 42.8% (zero imputation corrupts fusion), and late fusion 82.4% (graceful degradation by omission). The modality dropout training was critical: ablation showed that removing modality dropout (fixed full-modality training) reduced missing-2-modality performance retention from 86.8% to 54.2%. Per-modality importance (Shapley values) revealed task-specific patterns: drug response dominated by chemical structure (58.4%) with expression contributing 41.6%; COVID dominated by clinical labs (42.8%) with CT imaging 34.6% and notes 22.6%; disease classification relied most heavily on MRI (48.2%) with EHR 32.8% and demographics 19.0%.

4.3 Foundation Model Contribution and Contrastive Pre-Training

Foundation model ablation results are presented in Table 3 and Figure 3. Replacing foundation model encoders with randomly initialised equivalents (matched architecture) reduced CMBFN performance: drug response AUROC dropped from 0.868 to 0.812 (delta -0.056), disease classification F1 from 0.924 to 0.882 (-0.042), protein function F1 from 0.902 to 0.836 (-0.066), COVID AUROC from 0.912 to 0.872 (-0.040). Foundation models contributed an average of +0.051 performance across tasks. Contrastive pre-training (on paired biomedical data before task-specific training) contributed an additional +0.024 on average: disease classification F1 0.924 with pre-training vs 0.904 without; COVID AUROC 0.912 vs 0.886; protein function F1 0.902 vs 0.872. Pre-training was most valuable for low-data tasks (COVID with 12,400 patients) and least valuable for large-data tasks (disease classification with 42,800), consistent with transfer learning expectations.

4.4 Multi-Parameter Comparison

Comprehensive comparison is presented in Figure 4. CMBFN dominated across prediction accuracy, missing modality robustness, and interpretability while maintaining acceptable computational cost. MuMoT provided the best non-FM alternative. MMVAE offered the best missing modality handling among baselines. Concatenation was computationally cheapest but

performance-limited. Late fusion provided interpretability through per-modality predictions but missed cross-modal interactions.

Table 3. Multi-modal fusion performance: accuracy across four tasks, missing modality robustness, and foundation model contribution

Metric	Con cat	Late Fusio n	CC A	MM VAE	Mu MoT	CMBF N
Drug resp. AUROC	0.762	0.788	0.776	0.806	0.832	0.868*
Disease F1	0.786	0.826	0.798	0.848	0.884	0.924*
Protein GO F1	0.808	0.838	0.812	0.856	0.876	0.902*
COVID AUROC	0.878	0.882	0.858	0.894	0.898	0.912*
1 miss. retain (%)	72.4	88.6	78.4	86.4	74.8	94.2*
2 miss. retain (%)	42.8	82.4	N/A	76.4	58.6	86.8*
FM contri bution	N/A	N/A	N/A	N/A	N/A	+0.051 avg

*Best in row ($p < 0.05$). Retain %: percentage of full-modality performance retained with missing modalities. FM contribution: average performance gain from using foundation model encoders vs random initialisation (CMBFN only; baselines do not use FMs). CCA limited to 2-modality pairs. Results averaged across 5-fold CV.




Figure 1. Performance across four multi-modal biology tasks (best metric per task)




Figure 2. Missing modality robustness: % of full-modality performance retained




Figure 3. Foundation model contribution (delta-AUROC/F1) in CMBFN vs random init




Figure 4. Multi-parameter multi-modal fusion comparison (normalised 0-1)

5. Discussion

This study establishes that foundation model-enabled cross-modal fusion with contrastive pre-training and modality dropout provides the most effective framework for integrating heterogeneous biological data.

5.1 The Foundation Model Advantage

CMBFN's +0.051 average performance gain from foundation model encoders quantifies the value of pre-trained domain-specific representations.

Foundation models trained on massive unlabelled biomedical data (ESM-2 on 250M protein sequences, BiomedCLIP on 15M image-text pairs, DNABERT-2 on 135 genomes, BioBERT on PubMed + PMC corpora) capture biological patterns that downstream task-specific training cannot learn from limited labelled data. This transfer learning advantage is largest for protein function (+0.066 F1), where ESM-2's evolutionary representations provide rich structural and functional priors unavailable to random initialisation.

5.2 Missing Modality: Clinical Reality

CMBFN's 86.8% performance retention with 2 missing modalities directly addresses the primary practical barrier to multi-modal clinical deployment. In real clinical settings, the 'ideal' patient with all 5+ modalities available is rare -- most patients have subsets based on the tests their providers ordered. CMBFN's ability to maintain near-optimal performance with partial data enables deployment in real clinical workflows rather than requiring comprehensive multi-modal profiling. The modality dropout training -- randomly removing modalities during training -- is the key architectural choice: it forces the model to learn robust cross-modal representations that do not rely on any single modality being present.

5.3 Limitations

This study has limitations. The foundation models used (ESM-2, BiomedCLIP, DNABERT-2, BioBERT) have hundreds of millions of parameters, requiring GPU inference that limits point-of-care deployment. Contrastive pre-training requires paired multi-modal data that is rare for many biology-specific modality combinations. The four benchmark tasks, while diverse, do not cover all important multi-modal biology scenarios (single-cell multi-omics, spatial transcriptomics + imaging, longitudinal wearable data + EHR). The UK Biobank population is predominantly European (88%); generalisability to other populations requires validation. The CMBFN architecture has many hyperparameters (modality dropout rate, foundation model choice, contrastive temperature) that may need task-specific tuning. Privacy-preserving variants (federated learning with differential privacy) were not evaluated but are essential for clinical deployment. Interpretability at the fused representation level remains challenging; per-modality Shapley values provide some insight but do not fully explain cross-modal interactions.

6. Conclusion

This systematic benchmarking of six multi-modal fusion approaches across four diverse biological tasks establishes CMBFN -- a Cross-Modal Biological Fusion Network combining modality-specific foundation model encoders (ESM-2, BiomedCLIP, DNABERT-2, BioBERT), cross-modal attention with modality dropout, and contrastive pre-training on paired biomedical data -- as the most effective framework for biological data integration. CMBFN achieved the highest performance on all four tasks: drug response AUROC 0.868 (vs MuMoT 0.832; $p < 0.001$), disease classification F1 0.924 (vs MuMoT 0.884; $p < 0.001$), protein function F1 0.902, and COVID outcome AUROC 0.912. Foundation model encoders contributed +0.051 average performance; contrastive pre-training added +0.024 more; and modality dropout enabled 86.8% performance retention with 2 missing modalities (vs 58.6% for MuMoT without missing-modality training). Multi-modal fusion provided +0.060 average improvement over best single modality, confirming substantial value of cross-modal integration. CMBFN code (PyTorch, integrating ESM-2, BiomedCLIP, DNABERT-2, BioBERT), pre-trained contrastive weights, and task-specific fine-tuning scripts are available at <https://github.com/WEDSU-MultiModal/CMBFN-2024> under Apache 2.0 licence. Future work should extend to single-cell multi-modal integration (CITE-seq, spatial transcriptomics), develop privacy-preserving federated CMBFN for multi-institutional clinical deployment, incorporate temporal multi-modal data (longitudinal EHR + wearable), and pursue prospective clinical validation in diagnostic workflows.

References

- Acosta, J.N., Falcone, G.J., Rajpurkar, P. and Topol, E.J. (2022) Multimodal biomedical AI. *Nature Medicine*, 28(9), pp. 1773-1784.
- Akbari, H., Yuan, L., Qian, R., Chuang, W.H., Chang, S.F. and Cui, Y. (2021) VATT: transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34, pp. 24206-24221.
- Andrew, G., Arora, R., Bilmes, J. and Livescu, K. (2013) Deep canonical correlation analysis. In: *Proceedings of the 30th International Conference on Machine Learning*, pp. 1247-1255.
- Baltrusaitis, T., Ahuja, C. and Morency, L.P. (2019) Multimodal machine learning: a survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2), pp. 423-443.
- Chen, R.J., Lu, M.Y., Williamson, D.F.K., Chen, T.Y., Lipkova, J. and Noor, Z. (2022) Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell*, 40(8), pp. 865-878.
- Ji, Y., Zhou, Z., Liu, H. and Davuluri, R.V. (2021) DNABERT: pre-trained bidirectional encoder representations from transformers model for DNA-language in genome. *Bioinformatics*, 37(15), pp. 2112-2120.
- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S. and So, C.H. (2020) BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4), pp. 1234-1240.
- Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z. and Lu, W. (2023) Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637), pp. 1123-1130.
- Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J. and Topol, E.J. (2023) Foundation models for generalist medical artificial intelligence. *Nature*, 616(7956), pp. 259-265.
- Ouyang, D., He, B., Ghorbani, A., Lungren, M.P., Ashley, E.A. and Liang, D.H. (2023) Video-based AI for beat-to-beat assessment of cardiac function. *Nature*, 580(7802), pp. 252-256.
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G. and Agarwal, S. (2021) Learning transferable visual models from natural language supervision. In: *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748-8763.
- Sharifi-Noghabi, H., Zuo, W., Peng, S. and Ester, M. (2019) MOLI: multi-omics late integration with deep neural networks for drug response prediction. *Bioinformatics*, 35(14), pp. i501-i509.
- Tsai, Y.H.H., Bai, S., Liang, P.P., Kolter, J.Z., Morency, L.P. and Salakhutdinov, R. (2019) Multimodal transformer for unaligned multimodal language sequences. In: *Proceedings of the 57th Annual Meeting of the ACL*, pp. 6558-6569.
- Wu, M. and Goodman, N. (2018) Multimodal generative models for scalable weakly-supervised learning. *Advances in Neural Information Processing Systems*, 31.
- Zhang, S., Xu, Y., Usuyama, N., Bagga, J., Tinn, R. and Preston, S. (2023) Large-scale domain-specific pretraining for biomedical vision-language processing. *arXiv preprint arXiv:2303.00915*.
- Krishnan, R., Rajpurkar, P. and Topol, E.J. (2022) Self-supervised learning in medicine and healthcare. *Nature Biomedical Engineering*, 6(12), pp. 1346-1352.
- Huang, S.C., Pareek, A., Seyyedi, S., Banerjee, I. and 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(1), p. 136.

- Rajendran, S., Xu, Z., Pan, W., Ghosh, A. and Wang, F. (2024) Data heterogeneity in federated learning with Electronic Health Records: case studies of risk prediction for acute kidney injury and sepsis diseases in critical care. *PLoS Digital Health*, 3(3), p. e0000117.
- Zhou, Y., Chia, M.A., Wagner, S.K., Ayhan, M.S., Williamson, D.J. and Struyven, R.R. (2023) A foundation model for generalizable disease detection from retinal images. *Nature*, 622(7981), pp. 156-163.
- Boecking, B., Usuyama, N., Bannur, S., Castro, D.C., Schwaighofer, A. and Hyland, S. (2022) Making the most of text semantics to improve biomedical vision-language processing. In: *Proceedings of the European Conference on Computer Vision*, pp. 1-21.
- Hotelling, H. (1936) Relations between two sets of variates. *Biometrika*, 28(3-4), pp. 321-377.

Declarations

Funding

This study was supported by the Spanish Ministry of Science and Innovation grant PID2023-152486OB-I00 (N. Bianchi). Computational resources provided by the Barcelona Supercomputing Centre (BSC-CNS) under grant BCV-2024-3-0062. The funding body had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

N. Bianchi is a co-inventor on a pending patent application (EP24202486.2) covering the CMBFN cross-modal biological fusion architecture. No other competing interests are declared.

Data Availability Statement

CMBFN code, pre-trained contrastive weights, and task-specific fine-tuning scripts are available at <https://github.com/WEDSU-MultiModal/CMBFN-2024> under Apache 2.0 licence. UK Biobank data via ukbiobank.ac.uk (application 86024). GDSC/CCLE from depmap.org. UniProt/AlphaFold2 from uniprot.org/alphafold.ebi.ac.uk. COVID cohort aggregated data available upon reasonable request.

Ethical Approval

UK Biobank data access under approved application 86024. COVID-19 cohort data use approved by participating institutional review boards (aggregate protocol WEDSU-IRB-2024-0186). All data de-identified. No additional ethics approval required for this secondary analysis.

Appendix A

CMBFN Architecture and Contrastive Pre-Training

This appendix provides the CMBFN architecture, foundation model integration, cross-modal attention mechanism, and contrastive pre-training protocol.

A.1 Foundation Model Integration

1a. ESM-2 (650M, frozen): protein sequences -> mean-pooled 1,280-d embedding -> linear projection to 512-d modality token

1b. BiomedCLIP (ViT-B/16, fine-tuned): medical images -> 512-d CLS token; fine-tuning with LoRA (rank 16) for task adaptation

2a. DNABERT-2 (117M, frozen): DNA sequences -> BPE tokenisation -> mean-pooled 768-d embedding -> linear projection 512-d

2b. BioBERT (110M, fine-tuned): clinical text -> [CLS] token embedding 768-d -> linear projection 512-d; LoRA fine-tuning

A.2 Contrastive Pre-Training

3a. Pre-training data: UK Biobank (42K imaging + EHR + genomics), CPTAC (11K omics pairs), MIMIC-IV (48K imaging + text + tabular)

3b. Loss: symmetric InfoNCE across all modality pairs; temperature 0.07; batch size 512; 100 epochs on 4 x A100

4a. Modality dropout: each modality dropped with $p=0.3$ per sample during training; forces robust cross-modal representations

4b. Fine-tuning: task-specific cross-modal attention (4 layers, 8 heads, $d=512$); uncertainty-aware late fusion with learned modality confidence