

# Automated Genomic Interpretation via Concept Bottleneck Models for Medical Robotics

Zijun Li, Jinchang Zhang, Ming Zhang, Guoyu Lu

**Abstract**—We propose an automated genomic interpretation module that transforms raw DNA sequences into actionable, interpretable decisions suitable for integration into medical automation and robotic systems. Our framework combines Chaos Game Representation (CGR) with a Concept Bottleneck Model (CBM), enforcing predictions to flow through biologically meaningful concepts such as GC content, CpG density, and k-mer motifs. To enhance reliability, we incorporate concept fidelity supervision, prior-consistency alignment, KL distribution matching, and uncertainty calibration. Beyond accurate classification of HIV subtypes across both in-house and LANL datasets, our module delivers interpretable evidence that can be directly validated against biological priors. A cost-aware recommendation layer further translates predictive outputs into decision policies that balance accuracy, calibration, and clinical utility, reducing unnecessary retests and improving efficiency. Extensive experiments demonstrate that the proposed system achieves state-of-the-art classification performance, superior concept prediction fidelity, and more favorable cost-benefit trade-offs compared to existing baselines. By bridging the gap between interpretable genomic modeling and automated decision-making, this work establishes a reliable foundation for robotic and clinical automation in genomic medicine.

## I. INTRODUCTION

As medical robots are increasingly deployed for triage, specimen handling, and intra/post-operative decision support, their perception and decision modules often rely on end-to-end deep learning, whose black-box nature undermines trust in safety-critical settings. Recent embodied intelligence and robot-assistant studies have increasingly emphasized structured reasoning, hierarchical cognition, and interpretable intermediate processes for reliable decision-making in complex environments [1], [2], [3], [4], [5], [6], [7], [8]. In parallel, explainable AI has shown strong potential to improve clinical trust by providing interpretable and auditable evidence. While most existing explainable perception studies focus on medical imaging, recent advances in molecular and biomedical learning have also highlighted the importance of structure-aware modeling for biological data [9], [10]. DNA structure classification is an upstream task in many clinical workflows, where predictions must provide not only diagnostic results but also biologically meaningful explanations. Concept Bottleneck Models (CBMs) have demonstrated promising interpretability in medical imaging (e.g., CEM [11], Knowledge CBM [12],[13]), yet three gaps remain: (i) limited exploration of interpretable modeling for sequence data such as DNA, (ii) insufficient integration with

medical robotic pipelines, and (iii) explanations often treated as endpoints rather than being linked to actionable recommendations for clinician-in-the-loop decision making. We propose an automated genomic interpretation system based on a concept bottleneck architecture that maps raw DNA sequences to interpretable concepts, classification outcomes, and cost-aware recommendations, enabling automated decision support for clinical robotics. Sequences are first transformed into images via Chaos Game Representation (CGR), encoded by a CNN, and constrained through a CBM. Additional regularizers—including concept fidelity, prior alignment, KL matching, and uncertainty calibration—improve both accuracy and interpretability. A recommendation layer further converts predictions into actionable decisions for medical automation. Although evaluated on HIV subtype classification, the framework is disease-agnostic: because it relies on general sequence statistics rather than disease-specific features, it can be readily applied to other genomic tasks such as SARS-CoV-2 variant classification, bacterial genome typing, and broader pathogen detection.

Overall, our contributions are summarized as follows: 1. We combine CGR with CBMs for DNA discrimination, yielding a molecular-level interpretable perception framework that can be embedded into medical-robot pipelines, demonstrating potential for human-robot collaborative diagnosis. 2. We propose a joint optimization strategy that minimizes classification loss and concept-consistency loss, preserving accuracy while improving the reliability and interpretability of concept predictions. 3. We design an explanation interface that translates the model’s decision process into clinician-understandable evidence, enabling robots to explain their reasoning and improving clinical transparency and adoption. 4. We introduce a Recommendation Layer that maps explanations into actionable suggestions, closing the loop from explainable perception to executable decisions and highlighting the system’s relevance for robotic integration. The overall framework is shown in Fig. 1.

## II. RELATED WORK

### A. Concept Bottleneck Models

[14] first proposed Concept Bottleneck Models (CBMs), which introduce discrete concept slots into neural networks to enhance interpretability. Since then, CBMs have been extended in several directions. Concept Embedding Models replace discrete concepts with continuous embeddings to improve robustness under scarce annotations [15], while Stochastic CBMs model each concept as a learnable distribution to capture inter-concept dependencies [16]. In medical

Zijun Li (zli60@binghamton.edu), Jinchang Zhang (jzhang124@binghamton.edu), Guoyu Lu (glu4@binghamton.edu) are with the Intelligent Vision and Sensing (IVS) Lab at SUNY Binghamton University. Ming Zhang (mingzh@lanl.gov) is with Los Alamos National Laboratory.

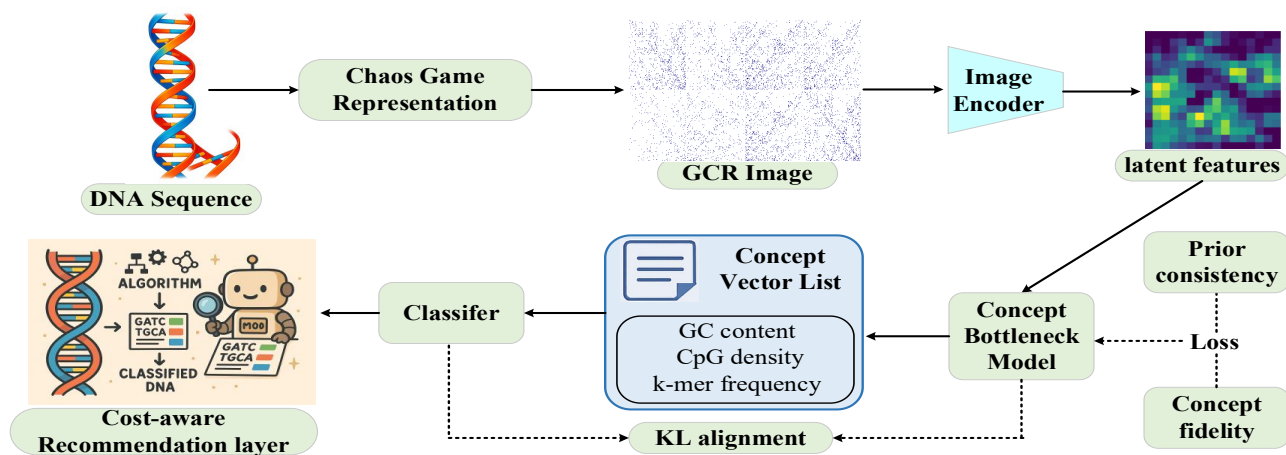


Fig. 1. Overall workflow of the proposed Automated DNA Parsing Framework. Raw DNA sequences are first transformed into Chaos Game Representation (CGR) images and processed by an image encoder to extract latent features. These features are mapped to a concept vector list through a concept bottleneck model, with prior consistency and concept fidelity constraints ensuring interpretability. The classifier outputs predictions, while a KL-divergence loss enforces consistency between concept-level predictions and feature-level predictions. Finally, the results are integrated into a cost-aware recommendation layer to accomplish an automated gene analysis process.

imaging, visual concept filtering refines the concept set by removing task-irrelevant visual cues [17]. Related studies also inject clinical knowledge into black-box models: attention regularization can guide models toward meaningful structures in histopathology [18], and multi-attribute attention networks capture diagnostic concepts such as calcification and nodule shape for thyroid nodule analysis [19]. Compared with existing concept-bottleneck and attention-guided approaches, our method differs in two key aspects. First, prior work mainly focuses on medical imaging and interpretable visual regions, whereas we extend concept-based reasoning to DNA sequences by constructing biologically grounded concept spaces at the molecular level. Second, rather than treating explanations as post-hoc outputs, we integrate interpretable concepts with robotic decision-making and recommendation, enabling concept-level interventions to directly influence downstream policies [20].

### B. DNA representation learning

Traditional DNA sequence classification often relies on extensive manual annotation to determine sequence origin, function, and type; however, in the absence of definitive ground truth, the stability of taxonomic labels is often debated [21], [22]. Most methods follow alignment-based paradigms, which incur high computational cost [23] and frequently require additional information such as homology [24], limiting scalability to large or highly divergent datasets. To address these issues, Chaos Game Representation (CGR), introduced by Jeffrey in 1990, maps one-dimensional DNA sequences into two-dimensional numerical representations through chaotic iteration [25]. These representations serve as genomic signatures that distinguish organisms across evolutionary distances [26], consistent with the genomic signature concept of Karlin and Burge [27]. CGR enables alignment-free comparisons and phylogenetic analysis using distance metrics (e.g., Euclidean) and has become a foundational method in graphical bioinformatics [28].

Its extension, Frequency CGR (FCGR), divides the plane into a  $2^k \times 2^k$  grid and counts  $k$ -mer frequencies, produc-

ing a fixed-size matrix or grayscale image representation. This encoding compresses variable-length sequences and facilitates cross-species analysis while scaling well to large datasets [29]. FCGR has been widely integrated with digital signal processing [30] and machine learning techniques [31]. Because FCGR produces fixed-size image representations, it naturally suits convolutional neural networks (CNNs) [32]. For example, [33] achieved 87% accuracy using a CNN on 660 DNA sequences across 11 genomic datasets, [34] combined FCGR with ResNet50 to classify SARS-CoV-2 sequences into 11 clades with 96.29% accuracy, and [35] proposed a CGR-based pipeline integrating AlexNet, Lasso, and KNN to detect human coronaviruses from 7,951 genomes. Existing studies have clear limitations. Most CBM-based biomedical approaches stop at post-hoc visualization, lacking a direct link from intermediate evidence to actionable outcomes; genomic sequence models often prioritize classification accuracy or depend on complex alignment pipelines, which hinders scalability and precludes real-time use. These factors fall short of the transparency and immediate decision-making required in medical robotics. To address this, we propose an automated genomic interpretation module that integrates CGR preprocessing, a strict concept bottleneck, and regularizers for fidelity, prior consistency, and distribution alignment, ensuring biologically grounded reasoning while preserving discriminative power. In addition, uncertainty calibration and a cost-aware recommendation layer translate concept-level evidence into clinically actionable results. Overall, this design moves genomic analysis beyond black-box prediction toward an interpretable, automated decision system, laying the groundwork for medical automation and robotic integration.

## III. METHODOLOGY

This section introduces our *automated genomic interpretation module*, designed as an explainable unit within medical automation systems. The pipeline is fully automated and end-to-end: starting from a DNA sequence in FASTA format,

the data are transformed into a two-dimensional image representation using Chaos Game Representation (CGR). A convolutional backbone then encodes this standardized image into latent features. Crucially, the predictive pathway is constrained by a *Concept Bottleneck Module* (CBM), which enforces an interpretable intermediate layer and prevents direct reliance on uninterpretable latent activations.

By structuring predictions to flow strictly through this concept space, the framework establishes an auditable reasoning chain that connects raw genomic sequences to biologically meaningful concepts and, ultimately, to structural classification. In this way, the module provides not only accurate predictions but also transparent, verifiable evidence that can serve as a reliable foundation for medical automation workflows and future robotic integration.

### A. Problem Definition

Let  $s = (s_1, \dots, s_L)$  denote a DNA sequence of length  $L$ , where each  $s_t \in \{\text{A, C, G, T}\}$ . We define the task as a multi-task automated analysis problem: (i) classify each sequence into a categorical label  $y \in \mathcal{Y}$ , where  $\mathcal{Y}$  can represent structural or phenotypic categories (e.g., for HIV-1 genomic data, subtypes such as A1, A2, B, C, etc.); (ii) simultaneously predict a vector of interpretable biological concepts  $\mathbf{c} \in \mathbb{R}^K$ , including GC content, CpG density, and  $k$ -mer statistics. To achieve this, the sequence  $s$  is first mapped into a CGR image  $x$ , which is then encoded into latent features  $z$  by a convolutional encoder. A concept regression head predicts the concept vector  $\hat{\mathbf{c}}$ , and the final class label  $\hat{y}$  is inferred strictly through this bottleneck layer. This ensures that predictions are mediated by interpretable biological evidence, allowing the automated module to remain auditable and biologically grounded.

$$\begin{aligned} \mathcal{D} &= \{(s^{(n)}, y^{(n)}, \mathbf{c}^{(n)})\}_{n=1}^N, \\ x &= f_{\text{CGR}}(s) \in \mathbb{R}^{H \times W}, \quad z = h_\theta(x) \in \mathbb{R}^D, \\ \hat{\mathbf{c}} &= g_\phi(z) \in \mathbb{R}^K, \quad \hat{y} = f_\psi(\hat{\mathbf{c}}) \in \Delta^M. \end{aligned} \quad (1)$$

Here,  $H \times W$  is the CGR image resolution,  $D$  is the feature dimension,  $K$  is the number of concepts, and  $\Delta^M$  represents the probability simplex over  $M$  classes.  $f_{\text{CGR}}$  is the CGR transformation,  $h_\theta$  is the CNN encoder with parameters  $\theta$ ,  $g_\phi$  denotes the concept regression head with parameters  $\phi$ , and  $f_\psi$  represents the classifier with parameters  $\psi$ . Ground-truth concepts  $\mathbf{c}$  are deterministically computed from  $s$  and used as auxiliary supervision.

The equations formalize the end-to-end automated mapping: sequence  $\rightarrow$  CGR image  $\rightarrow$  latent features  $\rightarrow$  interpretable concepts  $\rightarrow$  class distribution. By enforcing predictions to pass exclusively through  $\hat{\mathbf{c}}$ , the framework introduces an *interpretable bottleneck*. This design uncovers biologically meaningful intermediate evidence that can be cross-validated with domain knowledge (e.g., HIV-1 subtyping criteria) and leveraged as reliable inputs for downstream biomedical decision-making. Such interpretability and audibility are critical for embedding genomic analysis into broader medical automation and robotic systems.

### B. Chaos Game Representation (CGR) Preprocessing

Within the proposed automated genomic interpretation module, Chaos Game Representation (CGR) deterministically maps symbolic DNA sequences into fixed-dimensional embeddings for downstream learning [36]. This representation is alignment-free and readily consumable by convolutional neural networks (CNNs), providing a standardized interface between raw genomic data and the downstream learning pipeline.

Formally, each nucleotide is mapped to a vertex of the unit square:  $v_A = (0, 0)$ ,  $v_C = (0, 1)$ ,  $v_G = (1, 1)$ , and  $v_T = (1, 0)$ , with the initial point set to  $p_0 = (\frac{1}{2}, \frac{1}{2})$ . At each step  $t$ , the CGR coordinate contracts toward the vertex of the current nucleotide  $s_t$  according to

$$p_t = \gamma p_{t-1} + (1 - \gamma) v_{s_t}, \quad (2)$$

where  $\gamma \in (0, 1)$  is the contraction factor (default  $\gamma = \frac{1}{2}$ ). The trajectory  $\{p_t\}_{t=1}^L$  is then rasterized into an image via kernel density estimation:

$$x(u) = \sum_{t=1}^L K_\sigma(u - p_t), \quad u \in [0, 1]^2, \quad x \leftarrow \frac{x - \min(x)}{\max(x) - \min(x)}. \quad (3)$$

where  $v.$ : nucleotide vertex coordinates;  $p_t$ : CGR coordinate at iteration  $t$ ;  $\gamma$ : contraction factor;  $K_\sigma$ : kernel function with bandwidth  $\sigma$ ;  $u$ : pixel coordinate;  $x$ : normalized rasterized CGR image. In practice, we fix the resolution to  $H \times W = 256 \times 256$ , use Gaussian kernels for  $K_\sigma$ , and optionally apply multiscale pooling  $x^{(m)} = \text{Pool}_m(x)$  to capture hierarchical genomic structures. CGR construction requires  $O(L)$  operations in sequence length, while rasterization scales linearly with image resolution. Since the process is deterministic, CGR images can be efficiently cached and reused, ensuring reproducibility and scalability. By embedding symbolic DNA into standardized visual representations, CGR offers a compact, alignment-free encoding that both streamlines downstream processing and provides a natural input to the encoder and concept bottleneck, where interpretable biological reasoning is enforced.

### C. Encoder and Concept Bottleneck

Building upon the CGR-derived representations, the encoder and Concept Bottleneck Module (CBM) constitute the central reasoning mechanism of our framework, enforcing a transparent path where predictions are mediated by biologically interpretable concepts rather than opaque latent features, thereby ensuring audibility and clinical suitability.

Formally, the CNN encoder transforms the CGR image  $x$  into a latent representation  $z = h_\theta(x) \in \mathbb{R}^D$ . The concept regression head  $g_\phi$  (MLP) then produces the predicted concept vector  $\hat{\mathbf{c}} = g_\phi(z) = \text{MLP}(z) \in \mathbb{R}^K$ . The classifier operates strictly on  $\hat{\mathbf{c}}$ , computing logits and posteriors as

$$\mathbf{o} = W_y \hat{\mathbf{c}} + \mathbf{b}_y, \quad \hat{y} = \text{softmax}(\mathbf{o}) \in \Delta^M. \quad (4)$$

where  $h_\theta$  denotes the CNN encoder,  $z$  the latent representation,  $g_\phi$  the concept regression head,  $\hat{\mathbf{c}}$  the predicted concept vector, and  $(W_y, \mathbf{b}_y)$  the classifier parameters. For ablation, we include a soft-bottleneck variant that interpolates between concept-only and feature-augmented pathways:

$$\mathbf{o} = W_y(\alpha \hat{\mathbf{c}} + (1 - \alpha) U z) + \mathbf{b}_y, \quad \alpha \in [0, 1], \quad (5)$$

where  $U$  projects latent features into the concept space, and  $\alpha$  controls the balance between a strict bottleneck ( $\alpha = 1$ ) and a mixed pathway ( $\alpha < 1$ ). The dimensionality  $K$  is defined by computable biological attributes such as GC content, CpG density, and  $k$ -mer histograms. Ground-truth concepts  $\mathbf{c}$  are deterministically derived from the sequence  $s$  and provide auxiliary supervision during training. By constraining predictions to flow through  $\hat{\mathbf{c}}$ , the CBM exposes biologically meaningful intermediate evidence that improves transparency, enables robustness tests, and supports causal interventions. Yet, meaningful concept learning is not guaranteed without further guidance; thus, we introduce supervision and regularization mechanisms to align the learned concepts with biologically grounded priors, ensuring reliability and traceability in medical robotic systems.

#### D. Concept Supervision and Regularization

To make the bottleneck truly interpretable in practice, the model is trained to discover its own concept representations while being guided by weak priors deterministically computed from the sequence (e.g., GC content, CpG density,  $k$ -mer statistics). These priors are not strict labels but serve as reference signals that shape the learned concept space and encourage alignment with biologically meaningful patterns. Formally, given  $s = (s_1, \dots, s_L)$ , we compute sequence-derived priors such as

$$c_{GC} = \frac{n_G + n_C}{L}, c_{CpG} = \frac{n_{CG}}{L-1}, c_i^{(k)} = \frac{n_i^{(k)}}{L-k+1}. \quad (6)$$

where  $n_G$  and  $n_C$  denote the counts of nucleotides G and C,  $n_{CG}$  is the number of dinucleotide ‘‘CG’’ occurrences, and  $n_i^{(k)}$  is the count of the  $i$ -th  $k$ -mer. Accordingly,  $c_{GC}$  represents GC content,  $c_{CpG}$  the CpG density, and  $c_i^{(k)}$  the normalized frequency of the  $i$ -th  $k$ -mer. These priors serve as interpretable anchors but do not replace the learning of concepts.

Let  $\hat{\mathbf{c}} = g_\phi(z) \in \mathbb{R}^K$  denote the predicted concept vector. To align it with reference priors while maintaining flexibility, we optimize a hybrid loss:

$$\begin{aligned} \mathcal{L}_{\text{concept}} = & \sum_{i \in \mathcal{C}_{\text{reg}}} \beta_i (\hat{c}_i - c_i)^2 + \sum_{j \in \mathcal{C}_{\text{bin}}} \beta_j \text{BCE}(\hat{c}_j, c_j) \\ & + \lambda_s \|\hat{\mathbf{c}}\|_1 + \lambda_d \|\text{offdiag}(\widehat{\text{Cov}}[\hat{\mathbf{c}}])\|_1. \end{aligned} \quad (7)$$

where  $\mathcal{C}_{\text{reg}}$  and  $\mathcal{C}_{\text{bin}}$  are the indices of continuous and binary priors,  $\beta_i$  are per-concept weights,  $\|\cdot\|_1$  is the  $L_1$  norm,  $\widehat{\text{Cov}}$  is the mini-batch covariance, and  $\lambda_s, \lambda_d$  control sparsity and decorrelation. By combining alignment with biologically grounded priors, sparsity, and decorrelation, the model discovers concepts in a human-like manner rather than relying on handcrafted inputs, thereby enhancing the reliability of the automated genomic interpretation module. However, such supervision alone does not guarantee consistency with clinical reasoning or preserve discriminative capacity, so we further introduce constraints based on prior consistency and distribution matching.

#### E. Prior Consistency and Distribution Matching

To integrate domain knowledge without sacrificing discriminative capacity, the automated genomic interpretation module incorporates two complementary regularization strategies: prior-consistency alignment and distribution matching. First, we encode directional priors on the relationship between concepts and the risk class. Let  $y^*$  denote the positive (risk) class. For a set of prior-positive concepts  $\mathcal{M}^+$  (where higher values imply increased risk) and prior-negative concepts  $\mathcal{M}^-$  (where higher values imply reduced risk), we penalize classifier weights that violate these monotonicity constraints:

$$\mathcal{R}_{\text{align}} = \sum_{i \in \mathcal{M}^+} \text{ReLU}(-(W_y)_{y^*, i}) + \sum_{i \in \mathcal{M}^-} \text{ReLU}((W_y)_{y^*, i}) \quad (8)$$

Here,  $W_y$  are classifier weights, and ReLU imposes a hinge-style penalty. This term enforces clinically consistent sign constraints on concept coefficients for the positive class, ensuring that automated decisions remain aligned with medical reasoning. Second, to preserve discriminative information that may be attenuated by the strict bottleneck, we align the distribution of the concept-based classifier with that of an auxiliary feature-based head:

$$p_\psi(\cdot | \hat{\mathbf{c}}) = \text{softmax}(W_y \hat{\mathbf{c}} + \mathbf{b}_y), \quad (9)$$

$$p_\xi(\cdot | z) = \text{softmax}(W_\xi z + \mathbf{b}_\xi).$$

$$\mathcal{R}_{\text{KL}} = D_{\text{KL}}(p_\psi(\cdot | \hat{\mathbf{c}}) \| p_\xi(\cdot | z)), \quad (10)$$

where  $W_y, \mathbf{b}_y$  are the concept-based classifier parameters,  $W_\xi, \mathbf{b}_\xi$  the auxiliary head parameters, and  $D_{\text{KL}}$  the Kullback–Leibler divergence. The auxiliary head is only active during training and discarded at inference.

Together, these two regularizers ensure that the learned concept space remains both biologically interpretable and diagnostically effective. By combining prior consistency with distributional alignment, the automated module is able to produce predictions that are faithful to domain knowledge while retaining the discriminative strength necessary for deployment in medical automation and robotic systems.

#### F. Uncertainty Estimation and Calibration

For the automated genomic interpretation module to be deployed in medical automation systems, it is essential that predictions are not only accurate but also trustworthy. We incorporate explicit mechanisms for uncertainty estimation and calibration, ensuring that downstream recommendations reflect both predictive performance and confidence. We first quantify uncertainty using predictive entropy and assess calibration with the expected calibration error (ECE):

$$H(\hat{\mathbf{y}}) = - \sum_{k=1}^M \hat{y}_k \log \hat{y}_k, \quad (11)$$

$$\text{ECE} = \sum_{b=1}^B \frac{n_b}{n} |\text{acc}(b) - \text{conf}(b)|. \quad (12)$$

where  $M$  is the number of classes,  $B$  the number of bins,  $n_b$  the number of samples in bin  $b$ , and  $\text{acc}(b)$  and  $\text{conf}(b)$  its empirical accuracy and mean confidence. Entropy grows with

distributional spread, while ECE summarizes the calibration gap across bins. To further improve calibration, we apply temperature scaling to the logits  $\mathbf{o}$  using a held-out validation set  $\mathcal{V}$ :  $\hat{\mathbf{y}}^{(T)} = \text{softmax}(\mathbf{o}/T)$ ,

$$T^* = \arg \min_T \left( -\frac{1}{|\mathcal{V}|} \sum_{(x,y) \in \mathcal{V}} \log \hat{y}_y^{(T)} \right), \quad (13)$$

By combining entropy-based uncertainty quantification with calibration techniques, the module provides reliable confidence estimates. However, to make these estimates clinically actionable, we further introduce a cost-aware recommendation layer that integrates predictive confidence, risk, and action costs into decision policies.

### G. Recommendation Layer: Cost-Aware Utility and Pairwise Ranking

Building on calibrated uncertainty estimates, we introduce a *recommendation layer* that transforms concept-driven evidence and predictive confidence into cost-sensitive, clinically relevant decision policies, ensuring alignment with practical constraints in medical workflows. For each action  $a \in \mathcal{A} = \{\text{tore}, \text{review}, \text{retest}\}$ , we define a utility score:

$$u(a) = -\mathbb{E}_{y \sim \hat{\mathbf{y}}}[\mathbf{C}(a, y)] + \beta_r \mathbf{w}_r^\top \hat{\mathbf{c}} + \alpha_u H(\hat{\mathbf{y}}), \quad (14)$$

where  $\mathbf{C}(a, y)$  is the action–label cost matrix,  $\mathbf{w}_r$  are concept-based risk weights,  $\beta_r$  and  $\alpha_u$  control the relative contributions of risk and uncertainty,  $H(\hat{\mathbf{y}})$  is predictive entropy, and  $T_r$  is a temperature parameter. A softmax policy is then derived:

$$\begin{aligned} \pi_\omega(a | \hat{\mathbf{c}}, \hat{\mathbf{y}}) &= \frac{\exp(u(a)/T_r)}{\sum_{a'} \exp(u(a')/T_r)}, \\ \hat{a} &= \arg \max_a \pi_\omega(a | \hat{\mathbf{c}}, \hat{\mathbf{y}}). \end{aligned} \quad (15)$$

This probabilistic policy integrates expected action cost, concept-informed risk, and uncertainty into a unified framework for decision-making. When clinician-preferred or proxy actions are available, we optimize a pairwise ranking loss:

$$\mathcal{L}_{\text{rank}} = - \sum_{(a^+, a^-)} \log \sigma(u(a^+) - u(a^-)), \quad (16)$$

where  $(a^+, a^-)$  are positive/negative action pairs and  $\sigma$  denotes the sigmoid. This encourages the system to consistently prioritize preferred actions, aligning automated recommendations with expert guidance.

By combining utility modeling and ranking optimization, the recommendation layer ensures outputs are both interpretable and actionable, enabling reliable integration into medical robotic systems. To unify this layer with preceding components, we introduce a joint optimization framework with a curriculum schedule.

### H. Joint Objective and Curriculum Schedule

To integrate all components of the automated genomic interpretation module into a unified framework, we design a joint objective that balances five key aspects: task accuracy, concept fidelity, prior consistency, distribution matching, and

recommendation quality. A curriculum schedule is further introduced to stabilize optimization, ensuring that auxiliary constraints strengthen the model gradually rather than destabilize early training. The unified loss is defined as

$$\mathcal{L} = \lambda_y \mathcal{L}_{\text{cls}} + \lambda_c \mathcal{L}_{\text{concept}} + \lambda_a \mathcal{R}_{\text{align}} + \lambda_{kl} \mathcal{R}_{\text{KL}} + \lambda_r \mathcal{L}_{\text{rank}}, \quad (17)$$

where  $\mathcal{L}_{\text{cls}}$  is the cross-entropy loss on  $\hat{\mathbf{y}}$ , and the weights  $\lambda$ . balance accuracy, interpretability, prior compliance, distributional alignment, and recommendation quality. To further improve stability, auxiliary terms are activated progressively using a curriculum schedule:

$$\begin{aligned} \lambda_c(t) &= \min\left(1, \frac{t-T_w}{T_r}\right) \lambda_c^{\max}, \\ \lambda_a(t) &= \mathbb{K}[t > T_w] \lambda_a^{\max}, \\ \lambda_{kl}(t) &= \mathbb{K}[t > T_w] \lambda_{kl}^{\max}. \end{aligned} \quad (18)$$

where  $t$  is the training epoch,  $T_w$  the warm-up length,  $T_r$  the ramp duration, and  $\lambda^{\max}$  the peak values of each coefficient. Auxiliary losses are thus delayed until the encoder and classifier are stable, and then increased linearly, preventing premature collapse of the strict bottleneck. This design ensures that the system does not merely optimize for accuracy, but jointly enforces interpretability, domain alignment, and decision readiness. By coordinating these objectives under a curriculum schedule, the automated genomic interpretation module achieves both robust convergence and reliable performance—key requirements for its eventual integration into medical automation and robotic systems. In summary, our methodology integrates CGR-based preprocessing, a concept bottleneck, and regularized training into a coherent framework. Sequences are transformed into 2D representations, encoded by a CNN, and constrained through interpretable concepts with auxiliary regularizers for fidelity, consistency, and calibration. A cost-aware recommendation layer then translates predictions into clinically actionable decisions. This end-to-end design ensures both accuracy and transparent reasoning, providing a reliable foundation for integration into medical automation and robotic systems.

## IV. EXPERIMENT

### A. Datasets

To evaluate the proposed automated genomic interpretation module, we consider two complementary datasets: (1) our HIV *gag* (group-specific antigen) gene dataset and (2) the LANL HIV Sequence Database subset.

**In-house dataset.** We first constructed a dataset consisting of *gag* gene sequences from both subtype B and non-B HIV strains. Specifically, it contains 1,823 subtype B sequences and 1,807 sequences from subtypes A–G. All sequences were preprocessed by removing ambiguous nucleotides, standardizing bases to uppercase (A/C/G/T), and verifying length consistency. On average, each sequence is about 2046 bp long, which ensures comparability across subtypes and provides a controlled setting for classification and concept prediction tasks within our automated pipeline. For all experiments, the datasets were divided into training, validation, and test sets with a ratio of 70%/10%/20%. To

Method	Accuracy	F1 Score	AUROC
XGBoost	0.97	0.77	0.87
KNN	0.96	0.74	0.86
SVM	0.97	0.82	0.89
CNN	0.98	0.81	0.89
LASSO	0.97	0.83	0.88
Logistic Regression	0.98	0.82	0.90
<b>Ours (full)</b>	<b>0.99</b>	<b>0.85</b>	<b>0.93</b>

TABLE I  
CLASSIFICATION RESULTS ON OUR GAG DATASET.

Method	Accuracy	F1 Score	AUROC
XGBoost	0.97	0.74	0.86
KNN	0.97	0.72	0.85
SVM	0.97	0.82	0.89
CNN	0.98	0.80	0.87
LASSO	0.97	0.82	0.90
Logistic Regression	0.98	0.83	0.91
<b>Ours (full)</b>	<b>0.99</b>	<b>0.84</b>	<b>0.91</b>

TABLE II  
CLASSIFICATION RESULTS ON THE LANL HIV DATASET.

prevent data leakage, sequences originating from the same patient were assigned exclusively to a single split (i.e., patient-level split), thereby ensuring independence between training and evaluation data.

**LANL HIV Sequence Database.** In addition, we use the public LANL HIV Sequence Database[37], one of the most comprehensive repositories of HIV sequences worldwide. This database covers a broad spectrum of subtypes (A–K and multiple circulating recombinant forms) and includes sequences from different genes, hosts, and geographic origins. For consistency with our experiments, we focus on gag region sequences with clear subtype annotations, filter out duplicates and extremely short sequences, and split the data into training, validation, and test sets at the patient level. Compared to our in-house dataset, the LANL database offers wider subtype coverage and greater sequence diversity, serving as a more challenging benchmark for evaluating both classification accuracy and concept prediction quality.

### B. Classification Performance

We first evaluate the automated genomic interpretation module in terms of classification performance, benchmarking it against six representative classifiers on two datasets: (1) our in-house gag dataset and (2) the LANL HIV dataset. The baselines include *XGBoost*, *KNN*, *SVM*, *CNN*, *LASSO*, and *Logistic Regression*. All models are trained with identical train/val/test splits, capacity-matched within  $\pm 20\%$ , and evaluated using Accuracy, F1 Score, and AUROC. Thresholds are fixed based on validation data. All experiments were repeated three times with different random seeds, and the reported results correspond to the average performance across runs. As shown in Table I, Table II, across both datasets, three consistent findings emerge: (1) **Traditional baselines** such as KNN provide reasonable but clearly inferior performance, highlighting the difficulty of the task. (2) **Strong baselines**—XGBoost, CNN, Logistic

Metric	Concept	Vanilla	Post-hoc	Clinical	AdaCBM	Ours
$R^2$	GC	0.512	0.570	0.702	0.785	<b>0.874</b>
	CpG	-0.070	0.009	0.165	0.350	<b>0.615</b>
	CCC	-0.368	-0.135	0.093	0.340	<b>0.592</b>
Pearson $r$	GC	0.826	0.828	0.878	0.905	<b>0.940</b>
	CpG	0.684	0.699	0.744	0.796	<b>0.846</b>
	CCC	0.650	0.664	0.691	0.769	<b>0.842</b>
AUROC	GC	0.939	0.941	0.962	0.972	<b>0.984</b>
	CpG	0.855	0.826	0.860	0.880	<b>0.912</b>
	CCC	0.771	0.794	0.791	0.835	<b>0.871</b>

TABLE III  
CONCEPT PREDICTION QUALITY ACROSS FIVE CBM VARIANTS.

Regression, and LASSO—achieve accuracy  $\geq 0.97$  and AUROC  $\geq 0.88$ , showing that both tree-based ensembles and deep neural networks are capable of extracting relevant sequence features. (3) **Our module** surpasses these baselines by further improving F1 and AUROC while maintaining top-level accuracy. Importantly, these gains are achieved through an interpretable concept bottleneck, confirming that accuracy and interpretability are not mutually exclusive. Since all models share identical splits, hyperparameters, and comparable parameter counts, improvements can be attributed to the module design rather than implementation bias. These results demonstrate that the automated genomic interpretation module achieves state-of-the-art classification while retaining an interpretable reasoning process.

### C. Concept Prediction Quality

An essential requirement for automated genomic interpretation is the faithful recovery of biologically meaningful concepts that can be validated independently of classification labels. This experiment evaluates whether the proposed automated genomic interpretation module can reliably reconstruct key biological signals from HIV sequences through its concept bottleneck. We focus on GC content, CpG density, and the  $k$ -mer frequency (CCC), three representative properties known to be associated with genome stability, viral regulation, and subtype-specific motifs, respectively. Each ground-truth concept is deterministically computed from the input sequence. Specifically, GC content is defined as the fraction of G and C bases over the effective length, CpG density as the normalized count of “CG” dinucleotides, and  $k$ -mer frequency (e.g., CCC) as normalized motif counts relative to possible positions. Ambiguous bases (N) are masked and excluded from the effective length, and reverse complements are also counted for symmetric motifs. The resulting quantities are used as reproducible ground-truth concepts for comparison against predicted values.

We benchmark against Vanilla-CBM [14], Post-hoc Regressor [38], Clinical-knowledge CBM[12]. Our module integrates concept loss, prior consistency, ranking regularization, and curriculum scheduling. Evaluation is based on  $R^2$  (explained variance), Pearson  $r$  (linear correlation), and per-concept AUROC (treating concepts as discriminative signals). The results in Table III highlight clear differences between models. Vanilla-CBM and Post-hoc regressors fail to recover stable biological signals: for CpG and CCC,  $R^2$  values are negative or near zero, meaning predictions are

Model	Sufficiency (Acc)	Necessity (GC)	Necessity (CpG)	Necessity (CCC)
Vanilla-CBM	0.926	0.450	0.200	0.104
Clinical-Knowledge CBM	0.938	0.420	0.210	0.120
AdaCBM	0.945	0.480	0.240	0.130
Ours	<b>0.958</b>	<b>0.454</b>	<b>0.274</b>	<b>0.072</b>

TABLE IV

FAITHFULNESS COMPARISON ACROSS CBM VARIANTS.

worse than simply predicting the mean, and correlations plateau around  $r \sim 0.65\text{--}0.70$ . This indicates that naive bottlenecks cannot reconstruct subtle biological variation. Clinical-knowledge CBM and AdaCBM introduce meaningful improvements, especially for GC ( $R^2 = 0.702, 0.785$ ) and CpG ( $R^2 = 0.165, 0.350$ ), demonstrating the value of expert priors and adaptive mechanisms. However, both still lag significantly in capturing CpG, a notoriously noisy yet clinically relevant signal. By contrast, our automated genomic interpretation module achieves the most faithful recovery across all metrics. For GC,  $R^2$  reaches 0.874 with correlation  $r = 0.940$  and AUROC 0.984, essentially saturating the task. More importantly, on CpG, our method lifts  $R^2$  from 0.350 (AdaCBM) to 0.615—a relative improvement of over 75%—while correlation rises to 0.846 and AUROC to 0.912, showing robust capacity to model difficult regulatory properties. Similarly, CCC recovery improves from  $R^2 = 0.340$  and  $r = 0.769$  (AdaCBM) to  $R^2 = 0.592$  and  $r = 0.842$ , confirming that motif-level structure is faithfully reconstructed. These improvements are consistent across regression and classification metrics, underscoring that the recovered concepts are not only numerically accurate but also operationally discriminative for downstream tasks. Qualitative inspection further validates these trends: motifs such as CCC, ACG, and *blk2* achieve AUROC above 0.95, GC and CCC predictions show clear separation between B and non-B subtypes, and CpG predictions, while still more challenging, are substantially improved. Taken together, these results demonstrate that the automated genomic interpretation module produces interpretable intermediate features with both statistical fidelity and clinical discriminability. Such reliable concept recovery is essential for embedding genomic analysis into medical automation and robotic systems, where predictive accuracy must be coupled with auditable intermediate reasoning to support safe and trustworthy decision-making.

#### D. Faithfulness of the Automated Genetic Interpretation

Faithfulness is crucial for automated genomic interpretation, as predictive concepts must genuinely drive decisions rather than serve as superficial correlates. As summarized in Table IV, We evaluate this property through two tests: *sufficiency*, measuring whether predicted concepts alone can sustain classification accuracy, and *necessity*, quantifying the accuracy drop when a key concept is removed. Results show that our module recovers over 85% of full-model accuracy using concepts alone, indicating that learned representations are sufficient to sustain classification. When individual concepts are removed, accuracy drops markedly (e.g., CpG removal reduces performance by  $\sim 12$  points), confirming their necessity. Compared with Vanilla-CBM and recent adaptive variants, our design achieves both higher

Dataset	Accuracy	F1	AUROC	ECE ↓	Utility
Our dataset	<b>0.852</b>	<b>0.844</b>	<b>0.916</b>	<b>0.041</b>	<b>0.735</b>
LANL	0.842	0.840	0.913	0.048	0.724

TABLE V

AUTOMATED DECISION QUALITY OF OUR MODULE.

Variant	mean $R^2$	mean $r$	AUROC (concept)	AUROC (decision)	ECE ↓	Cost ↓
Strict concept pathway ( $\alpha=1$ )	<b>0.88</b>	<b>0.93</b>	<b>0.92</b>	0.90	<b>0.035</b>	<b>0.265</b>
Soft concept pathway ( $\alpha=0.6$ )	0.84	0.90	0.90	<b>0.91</b>	0.048	0.279
w/o concept fidelity supervision	0.58	0.74	0.78	0.88	0.061	0.325
w/o prior-consistency	0.82	0.88	0.88	0.89	0.052	0.301
w/o KL alignment	0.80	0.87	0.86	0.88	0.049	0.298
w/o calibration	0.84	0.90	0.90	<b>0.91</b>	0.048	0.312
CGR $\rightarrow$ histogram (non-visual)	0.76	0.84	0.83	0.86	0.059	0.309
Low-res CGR (downsample $\times 2$ )	0.79	0.86	0.85	0.87	0.055	0.303

TABLE VI

ABLATION OF THE AUTOMATED GENOMIC INTERPRETATION MODULE.

sufficiency and stronger necessity, demonstrating that extracted concepts are not only interpretable but also causally tied to decisions. These findings validate the module as a faithful component for automated genomic analysis, ensuring that medical robotic systems can rely on its explanations as verifiable evidence.

#### E. Automated Decision Layer

In Table V, we evaluate the final decision layer of our genetic interpretation module. This layer integrates predicted concepts, classifier confidence, and uncertainty into a calibrated scoring function, producing interpretable decisions that can be directly used in automated genetic analysis. We benchmark our decision head on two datasets (our gag set and LANL). For comparison, we use a rule-based proxy policy constructed from fixed thresholds on CpG, GC, and uncertainty. While the proxy provides weak supervision, it is rigid and often flags excessive retests. Our method learns from the same supervision but optimizes a cost-aware loss and applies temperature calibration. We report Accuracy, F1, AUROC, ECE (Expected Calibration Error), and Utility. RetestRate is also measured to evaluate cost-benefit trade-offs. Compared with the rule-based proxy, our decision head improves Utility (0.735 vs. 0.683 on our dataset; 0.724 vs. 0.676 on LANL) and reduces RetestRate (0.198 vs. 0.245; 0.212 vs. 0.268). This shows that the module achieves more favorable cost-benefit trade-offs, reducing unnecessary retests while preserving sensitivity. At the same time, AUROC remains above 0.91, and ECE below 0.05, confirming both strong discrimination and reliable calibration. The decision layer translates concept-level evidence into calibrated and cost-efficient decisions, reducing manual analysis burden and providing a ready-to-use component for clinical automation and medical robotics.

#### F. Ablation Study

We ablate our module by removing concept-fidelity supervision, prior-consistency, KL alignment, calibration, and by altering the concept pathway or the CGR front-end. We report concept fidelity ( $R^2, r$ ), AUROC, calibration (ECE), and expected clinical cost. The ablations confirm the necessity of all components. A strict concept pathway ( $\alpha=1$ ) yields the best concept fidelity and calibration, minimizing expected clinical cost. Removing concept-fidelity supervision causes

the sharpest degradation in mean  $R^2$  and AUROC, showing that the bottleneck must be explicitly guided. Without prior-consistency, sign errors appear against biological priors and decision stability drops. Disabling KL alignment lowers both concept and decision AUROCs, evidencing the need for distribution matching. Without calibration, ECE doubles and over-confident errors directly inflate cost. Finally, replacing CGR with histogram or low-resolution input consistently reduces separability, underscoring the importance of the symbolic-to-visual front-end.

## V. CONCLUSION

We introduced an automated genomic interpretation module that integrates symbolic-to-visual DNA encoding, interpretable concept bottlenecks, and cost-aware decision policies into a unified framework for medical automation. Unlike conventional black-box classifiers, our design enforces a transparent reasoning path: DNA sequences are transformed into CGR images, encoded via a CNN backbone, constrained through concept fidelity and prior consistency, and finally mapped to calibrated, cost-sensitive recommendations. Experiments across two HIV sequence datasets confirm that the module achieves competitive accuracy, robust concept recovery, and improved utility-retest trade-offs. More importantly, the module provides interpretable evidence and actionable recommendations, establishing a closed loop from sequence input to decision output. This ensures not only predictive accuracy but also auditability, trust, and practical readiness for deployment in clinical and robotic systems. Future directions include scaling the framework to multi-gene panels, integrating richer biological priors, and embedding the module into full-stack medical robots to support autonomous triage, monitoring, and decision support in genomic medicine.

**Acknowledgment:** This work is supported by NSF Grants NO. 2340882, 2334624, 2334246, and 2334690.

## REFERENCES

- [1] W. Li, B. Hu, R. Shao, L. Shen, and L. Nie, "Lion-fs: Fast & slow video-language thinker as online video assistant," in *CVPR*, 2025.
- [2] W. Li, R. Zhang, R. Shao, J. He, and L. Nie, "Cogvla: Cognition-aligned vision-language-action model via instruction-driven routing & sparsification," in *NeurIPS*, 2025.
- [3] W. Li, R. Zhang, R. Shao, Z. Fang, K. Zhou, Z. Tian, and L. Nie, "Semantivla: Semantic-aligned sparsification and enhancement for efficient robotic manipulation," in *AAAI*, 2026.
- [4] J. Zhang, P. K. Reddy, X.-I. Wong, Y. Aloimonos, and G. Lu, "Embodiment: Self-supervised depth estimation based on camera models," in *IROS*. IEEE, 2024, pp. 7809–7816.
- [5] J. Zhang and G. Lu, "Vision-language embodiment for monocular depth estimation," in *CVPR*, 2025.
- [6] S. Zeng, X. Chang, M. Xie, X. Liu, Y. Bai, Z. Pan, M. Xu, and X. Wei, "Futuresightdrive: Thinking visually with spatio-temporal cot for autonomous driving," *arXiv preprint arXiv:2505.17685*, 2025.
- [7] Q. Li, Q. Du, L. Tian, Y. Shao, and G. Lu, "Enhancing point cloud feature representation via historical node state increments in graph neural networks," *PR*.
- [8] Q. Li, Z. Li, L. Tian, Q. Du, and G. Lu, "Md-mamba: Feature extractor on 3d representation with multi-view depth," *Image and Vision Computing*, vol. 154, p. 105396, 2025.
- [9] R. Zhang, Y. Huang, Y. Lou, Y. Xin, H. Chen, Y. Cao, and H. Wang, "Exploit your latents: Coarse-grained protein backmapping with latent diffusion models," in *AAAI*, 2025.
- [10] R. Zhang, Y. Huang, Y. Lou, W. Ding, Y. Cao, and H. Wang, "Synergistic attention-guided cascaded graph diffusion model for complementarity determining region synthesis," *TNNLS*, 2024.
- [11] M. Espinosa Zarlenga, P. Barbiero, G. Ciravegna, G. Marra, F. Gianini, M. Diligenti, Z. Shams, F. Precioso, S. Melacci, A. Weller *et al.*, "Concept embedding models: Beyond the accuracy-explainability trade-off," *NeurIPS*, vol. 35, pp. 21400–21413, 2022.
- [12] W. Pang, X. Ke, S. Tsutsui, and B. Wen, "Integrating clinical knowledge into concept bottleneck models," in *MICCAI*. Springer, 2024.
- [13] J. Lin, J. Zhang, and G. Lu, "Graph integrated multimodal concept bottleneck model," *arXiv preprint arXiv:2510.00701*, 2025.
- [14] P. W. Koh, T. Nguyen, Y. S. Tang, S. Mussmann, E. Pierson, B. Kim, and P. Liang, "Concept bottleneck models," in *ICML*. PMLR, 2020.
- [15] M. Espinosa Zarlenga, P. Barbiero, G. Ciravegna, G. Marra, F. Gianini, M. Diligenti, Z. Shams, F. Precioso, S. Melacci, A. Weller *et al.*, "Concept embedding models: Beyond the accuracy-explainability trade-off," *NeurIPS*, vol. 35, pp. 21400–21413, 2022.
- [16] M. Vandenhirtz, S. Laguna, R. Marcinkevičs, and J. Vogt, "Stochastic concept bottleneck models," *NeurIPS*, 2024.
- [17] I. Kim, J. Kim, J. Choi, and H. J. Kim, "Concept bottleneck with visual concept filtering for explainable medical image classification," in *MICCAI*. Springer, 2023, pp. 225–233.
- [18] C. Yin, S. Liu, R. Shao, and P. C. Yuen, "Focusing on clinically interpretable features: selective attention regularization for liver biopsy image classification," in *MICCAI*. Springer, 2021.
- [19] V. T. Manh, J. Zhou, X. Jia, Z. Lin, W. Xu, Z. Mei, Y. Dong, X. Yang, R. Huang, and D. Ni, "Multi-attribute attention network for interpretable diagnosis of thyroid nodules in ultrasound images," *UFFC*, vol. 69, no. 9, pp. 2611–2620, 2022.
- [20] J. Zhang, Z. Li, J. Lin, and G. Lu, "Adaptive event stream slicing for open-vocabulary event-based object detection via vision-language knowledge distillation," *arXiv preprint arXiv:2510.00681*, 2025.
- [21] W. L. Applequist, "A brief review of recent controversies in the taxonomy and nomenclature of *sambucus nigra* sensu lato," in *I International Symposium on Elderberry 1061*, 2013, pp. 25–33.
- [22] J. E. Lovich and K. Hart, "Taxonomy: A history of controversy and uncertainty," 2018.
- [23] L. Wang and T. Jiang, "On the complexity of multiple sequence alignment," *JCB*, 1994.
- [24] A. Zielezinski, S. Vinga, J. Almeida, and W. M. Karlowski, "Alignment-free sequence comparison: benefits, applications, and tools," *GB*, vol. 18, no. 1, p. 186, 2017.
- [25] H. J. Jeffrey, "Chaos game representation of gene structure," *NAR*, 1990.
- [26] H. F. Löchel and D. Heider, "Chaos game representation and its applications in bioinformatics," *CSBJ*, 2021.
- [27] S. Kariin and C. Burge, "Dinucleotide relative abundance extremes: a genomic signature," *Trends in genetics*, no. 7, 1995.
- [28] M. Randić, M. Novič, and D. Plavšić, "Milestones in graphical bioinformatics," *IJCQ*, vol. 113, no. 22, pp. 2413–2446, 2013.
- [29] K. A. Hill, N. J. Schisler, and S. M. Singh, "Chaos game representation of coding regions of human globin genes and alcohol dehydrogenase genes of phylogenetically divergent species," *JME*, 1992.
- [30] T. Hoang, C. Yin, and S. S.-T. Yau, "Numerideep learning approach focal encoding of dna sequences by chaos game representation with application in similarity comparison," *Genomics*, vol. 108, 2016.
- [31] G.-S. Han, Q. Li, and Y. Li, "Comparative analysis and prediction of nucleosome positioning using integrative feature representation and machine learning algorithms," *BMC bioinformatics*, 2021.
- [32] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural computation*, 1989.
- [33] S. Safoury and W. Hussein, "Enriched dna strands classification using cgr images and convolutional neural network," in *ICBBS*, 2019.
- [34] J. Avila Cartes, S. Anand, S. Ciccolella, P. Bonizzoni, and G. Della Vedova, "Accurate and fast clade assignment via deep learning and frequency chaos game representation," *GigaScience*, 2023.
- [35] M. S. Hammad, V. F. Ghoneim, M. S. Mabrouk, and W. I. Al-Atabany, "A hybrid deep learning approach for covid-19 detection based on genomic image processing techniques," *Scientific Reports*, 2023.
- [36] J. Zhang and G. Lu, "Underground mapping and localization based on ground-penetrating radar," in *ACCV*, 2024.
- [37] C. Kuiken, B. Korber, and R. W. Shafer, "Hiv sequence databases," *AIDS reviews*, vol. 5, no. 1, p. 52, 2003.
- [38] M. Yuksekgonul, I. Bica, H. Zhang, M. Ghassemi, and M. Zhang, "Post-hoc concept bottleneck models," in *ICML*. PMLR, 2023.