

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 A RESOLUTION-AGNOSTIC GEOMETRIC TRANSFORMER FOR CHROMOSOME MODELING USING INERTIAL FRAME

006
007 **Anonymous authors**
008 Paper under double-blind review

011 ABSTRACT

013 Chromosomes are the carriers of genetic information. Further understanding their
014 3D structure can help reveal gene-regulatory mechanisms and cellular functions.
015 A standard pipeline for reconstructing the chromosome 3D structure first applies
016 the single-cell Hi-C high-throughput screening method to measure pairwise inter-
017 actions between DNA fragments at different resolutions; then it adopts computa-
018 tional methods to reconstruct the 3D structures from these contacts. These include
019 traditional numerical methods and deep learning models, which struggle with lim-
020 ited model expressiveness and poor generalization across resolutions. To solve
021 this issue, we propose InertialGenome, a novel transformer-based framework for
022 robust and resolution-agnostic chromosome reconstruction. InertialGenome first
023 adopts the inertial frame for the pose canonicalization. Then, based on such an
024 invariant frame, it proposes a Transformer with geometry-aware positional encod-
025 ing, leveraging Nyström estimation. To verify the effect of InertialGenome, we
026 evaluate our model on two single-cell 3D reconstruction datasets with four resolu-
027 tions, reaching superior performance over all four computational baselines. In ad-
028 dition to the structure metrics, we observe that InertialGenome outperforms when
029 analyzing the function of reconstructed structures on two validation tasks. Finally,
030 we leverage InertialGenome for cross-resolution transfer learning, yielding up to
031 a 5% improvement from low to high resolution.

032 1 INTRODUCTION

033 The genome encodes the complete set of genetic information within an organism, stored as a full
034 DNA sequence. This information is packaged into chromosomes, which serve as the carriers of ge-
035 netic material and compact the DNA into three-dimensional structures. These chromosomes adopt
036 complex 3D conformations that play essential roles in gene regulation, cell differentiation, and dis-
037 ease progression (Lieberman-Aiden et al., 2009; Dixon et al., 2012; Rao et al., 2014). Importantly,
038 such an organization cannot be inferred from the linear DNA sequence alone (Consortium et al.,
039 2024).

040 Over the past two decades, diverse experimental techniques have emerged to probe the 3D confor-
041 mations. Early approaches such as 3C(Dekker et al., 2002), 4C(Simonis et al., 2006), and 5C(Dostie
042 et al., 2006) enable targeted interrogation of chromatin interactions at specific loci. More recent
043 genome-wide methods including ChIA-PET(Fullwood et al., 2009), SPRITE(Quinodoz et al., 2022),
044 GAM(Beagrie et al., 2017), and Hi-C(Belton et al., 2012) provide comprehensive maps of spatial
045 chromatin contacts across the entire genome. Among them, high-throughput chromosome confor-
046 mation capture (Hi-C) enables genome-wide profiling of interactions between genomic loci. In
047 Hi-C, the genome is partitioned into consecutive, non-overlapping segments called *bins*, with their
048 length determined by the chosen *resolution* (e.g., 1 kb, 10 kb, or 250 kb per bin, with higher reso-
049 lution corresponding to shorter segments). As a result, Hi-C produces a contact matrix that records,
050 for every pair of bins, the frequency of spatial contacts between their underlying DNA segments,
051 as illustrated in Figure 1A. Hi-C maps at different resolutions offer complementary insights: high-
052 resolution maps capture fine local structures but are often sparse and noisy, while low-resolution
053 maps are denser and more robust, reflecting global organization. We focus on this *cross-resolution*

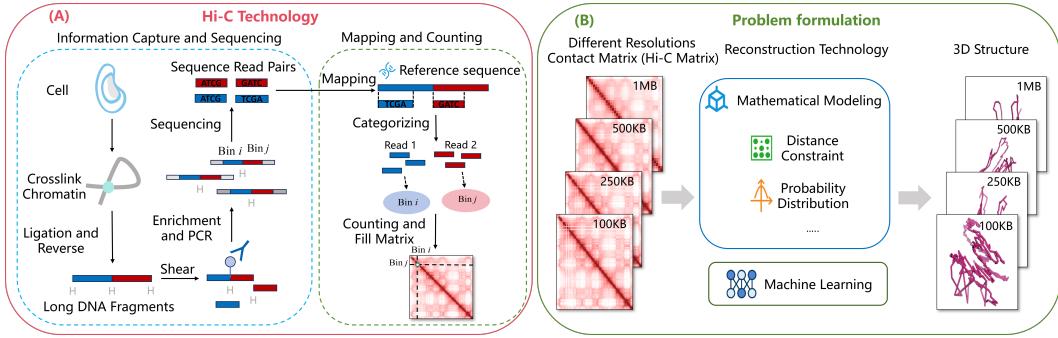


Figure 1: Overview of chromosome 3D reconstruction using Hi-C technology. (A): Experimental implementation of Hi-C for obtaining contact matrix information. (B): Computational pipeline for 3D structure reconstruction via mathematical modeling or machine learning based on the Hi-C contact matrix.

task by using low-resolution maps as structural priors to guide high-resolution 3D genome reconstruction.

Following the Hi-C matrix, the next step is to computationally reconstruct the 3D conformations (Trieu & Cheng, 2017; Oluwadare et al., 2018; Wang & Cheng, 2025). Traditional numerical methods based on distance geometry, such as ChromSDE (Zhang et al., 2013), miniMDS (Rieber & Mahony, 2017), and 3DMAX (Oluwadare et al., 2018), are computationally intensive and largely limited to low-resolution modeling due to data sparsity. To address these issues, deep learning offers an efficient, data-driven alternative. For instance, HiC-GNN (Hovenga et al., 2023) leverages node embeddings and graph neural networks to predict 3D conformation directly from Hi-C contact graphs; HiCEGNN (Wang & Cheng, 2025) further incorporates $E(3)$ -equivariance as a model constraint. However, these approaches share some key limitations: they rely solely on Hi-C contacts without incorporating explicit geometric priors (e.g., chromatin principal axes or directional chain structure). Moreover, the strong symmetry constraints of structures such as HiCEGNN limit model expression ability, making it difficult to process asymmetric structures (such as anchored loops).

Our Contributions. To tackle this problem, we propose InertialGenome, a novel Transformer-based framework for 3D chromosome reconstruction. InertialGenome has two main components. (1) InertialGenome performs pose canonicalization. It aligns each chromosome to its inertial frame—a coordinate system defined by the principal axes of its inertia tensor. This tensor is computed from the 3D point cloud of the chromosome. The alignment removes arbitrary rotations and translations, resulting in a pose-invariant representation. (2) Based on this invariant frame, InertialGenome employs a geometry-aware positional encoding into the Transformer architecture. The core idea is to project the bin-wise positions into an imaginary space and use the Nyström method to estimate the pairwise distance via inner products. Specifically, Nyström enables efficient low-rank estimation of the radial basis function (RBF) kernel over 3D coordinates, capturing long-range structural dependencies without computing the full distance matrix. To verify the effectiveness of InertialGenome, we conduct experiments on two single-cell 3D chromosome reconstruction datasets at four resolutions, where it consistently outperforms four baselines in both two structure metrics and two functional validation tasks. Additionally, InertialGenome excels at cross-resolution transfer tasks, reaching up to a 5% performance improvement.

Related Work. Existing methods for 3D chromosome reconstruction fall into three categories (see Appendix B): **Distance-based methods** (e.g., 3DMAX (Oluwadare et al., 2018), LorDG (Trieu & Cheng, 2017)) convert contacts to distance constraints; **Probabilistic approaches** (e.g., BACH (Hu et al., 2013), PASTIS (Varoquaux et al., 2014)) model contact matrices as observations from spatial distributions; **Deep learning methods** (e.g., HiC-GNN (Hovenga et al., 2023), HiCEGNN (Wang & Cheng, 2025)) map interactions to 3D structures via neural networks. However, these methods have some limitations: they rely on simplistic modeling of contact matrices as the sole input, lacking deeper structural interpretation, and their model expression ability may be constrained and limited. To overcome these issues, we propose InertialGenome, a Transformer-based framework that integrates inertial frame canonicalization and geometry-aware positional encoding, significantly improving robustness across resolutions.

108
109

2 PRELIMINARIES

110
111 The 3D chromosome reconstruction methods typically take the Hi-C contact matrix as input. The
112 resolution of a Hi-C contact matrix reflects the length of chromosome segments, with higher resolution
113 corresponding to shorter segments. For example, at 1 kb resolution, each 1,000 base pairs forms
114 a segment. When applied to the human genome with standard masking of unmappable regions, this
115 yields approximately 248,947 bins. This underscores the sharp scale differences across resolutions,
116 with high-resolution 3D reconstruction being far more computationally demanding and challenging.

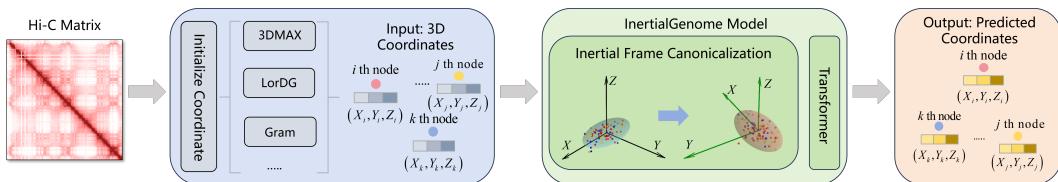
117
118 The reconstruction process consists of two key steps, as shown in Figure 1B. First, the contact
119 frequencies IF_{ij} are converted to spatial distances D_{ij} based on the inverse relationship between
120 distance and contact frequency (Pombo & Nicodemi, 2014; Barbieri et al., 2012), expressed as
121 $D_{ij} = IF_{ij}^{-\gamma}$ where $\gamma \in [0.1, 0.2, \dots, 2]$. Second, the 3D coordinates (x_i, y_i, z_i) are inferred from
122 these distances using either numerical methods or deep learning methods. In the following paper, for
123 each bin i , we denote by t_i the index of the i -th bin in the chromosome, and $\mathbf{s}_i = (s_{x_i}, s_{y_i}, s_{z_i}) \in \mathbb{R}^3$
124 its canonicalized 3D coordinate i.e., pose-normalized via alignment to the chromosome’s principal
125 axes (see Section 3.1).

126
127 **Problem formulation.** In this work, we are interested in solving the 3D chromosome reconstruction
128 task. Following the existing paradigm (Hovenga et al., 2023; Wang & Cheng, 2025), the first step
129 is to apply a numerical method to generate an initial position $\mathbf{C}^* \triangleq \{(x_i^*, y_i^*, z_i^*)\}_{i=1}^N$, where N is
130 the number of bins (e.g., nodes), from Hi-C contact matrix. Then, our model takes these initial 3D
131 coordinates \mathbf{C}^* as the input, and the output is an accurately reconstructed 3D coordinates $\hat{\mathbf{C}} \triangleq$
132 $\{(\hat{x}_i, \hat{y}_i, \hat{z}_i)\}_{i=1}^N$. Rigorously, we are solving the chromosome 3D reconstruction task as $\hat{\mathbf{C}} = f(\mathbf{C}^*)$.

133
134

3 METHOD INERTIALGENOME

135
136 In this section, we introduce InertialGenome, a novel Transformer-based framework for robust and
137 resolution-agnostic chromosome reconstruction. It consists of three key components: inertial frame
138 canonicalization, geometry-aware positional encoding, and structure-aware fusion. Figure 2 shows
139 the whole architecture of InertialGenome.



140
141
142
143
144 Figure 2: The architecture of InertialGenome. The model takes as input 3D chromosome coordinates
145 reconstructed by numerical methods. It first performs inertial-frame canonicalization to
146 obtain a pose-invariant representation. Then, based on this invariant frame, InertialGenome applies
147 a Transformer with geometry-aware positional encoding, leveraging Nyström estimation, to process
148 the data. The framework then outputs stable 3D chromosome structures.

149
150

3.1 INERTIAL FRAME CANONICALIZATION

151
152 To achieve pose-invariant representation of 3D chromosome structures, we implement an inertial
153 frame canonicalization method, with the following steps: (1) **Centroid translation**: $\bar{c} =$
154 $\frac{1}{N} \sum_{i=1}^N \mathbf{c}_i$, where $\mathbf{c}_i \in \mathbf{C}^*$. It will adjust position relative to the center $\mathbf{c}'_i = \mathbf{c}_i - \bar{c}$. (2) **Inertia tensor
155 computation**: We estimate the normalized inertia tensor as $\hat{\mathbf{I}} = \frac{1}{N} \sum_{i=1}^N (\|\mathbf{c}'_i\|^2 \mathbf{I}_3 - \mathbf{c}'_i (\mathbf{c}'_i)^T)$,
156 where \mathbf{I}_3 is the 3×3 identity matrix. (3) **Principal axes alignment**: We perform eigen-
157 decomposition of the inertia tensor $\hat{\mathbf{I}} = L \Lambda L^T$, where $\Lambda = \text{diag}(\lambda_x, \lambda_y, \lambda_z)$ contains the eigen-
158 values with $\lambda_x \geq \lambda_y \geq \lambda_z$. The columns of L are the corresponding orthonormal eigenvectors $\mathbf{l}_x, \mathbf{l}_y, \mathbf{l}_z$,
159 which define the principal axes in descending eigenvalue order. (4) **Chirality correction**: We select
160 the farthest point $\mathbf{c}_{\max} = \arg \max_i \|\mathbf{c}'_i\|$ and map it into the principal-axis frame: $\mathbf{p} = L^T \mathbf{c}_{\max}$,
161 where $\mathbf{p} = (p_x, p_y, p_z)$ is the coordinate of the farthest point expressed in the principal-axis basis.
162 We then adjust the first two axes as $\mathbf{l}_x \leftarrow \text{sign}(p_x) \mathbf{l}_x$, $\mathbf{l}_y \leftarrow \text{sign}(p_y) \mathbf{l}_y$, and enforce a right-handed

system by $\mathbf{l}_z = \mathbf{l}_x \times \mathbf{l}_y$. Then, the canonical transform is $R = [\mathbf{l}_x, \mathbf{l}_y, \mathbf{l}_z]^T$. Pose-invariant coordinates are obtained as $\mathbf{s}_i = R\mathbf{c}'_i$, where $\mathbf{s}_i \in \mathbf{S}$, $\mathbf{S} \triangleq \{\mathbf{s}_{x_i}, \mathbf{s}_{y_i}, \mathbf{s}_{z_i}\}_{i=1}^N$. Although chromosome structures naturally occupy 3D space, reconstruction algorithms may yield degenerate cases (e.g., nearly coplanar nodes). We discard such samples before training. After applying the canonical transform, each coordinate is represented as

$$\mathbf{s}_i = R\mathbf{c}'_i, \quad \mathbf{S} \triangleq \{\mathbf{s}_i\}_{i=1}^N \subset \mathbb{R}^3. \quad (1)$$

3.2 GEOMETRY-AWARE POSITIONAL ENCODING

In 3D chromosome structure modeling, we need a position embedding method that can maintain both absolute position information and the relative distance when calculating the inner product, *e.g.*, $R_{\mathbf{s}_{x_1}, \mathbf{s}_{y_1}, \mathbf{s}_{z_1}}^T R_{\mathbf{s}_{x_2}, \mathbf{s}_{y_2}, \mathbf{s}_{z_2}} = R_{\mathbf{s}_{x_1} - \mathbf{s}_{x_2}, \mathbf{s}_{y_1} - \mathbf{s}_{y_2}, \mathbf{s}_{z_1} - \mathbf{s}_{z_2}}$.

(1) Geometric Position Encoding with RoPE. Inspired by Su et al. (2024), we extend Rotary Position Embedding (RoPE) to 3D Euclidean space. To achieve rotation-equivariant attention while maintaining pairwise distance awareness, we decompose 3D spatial encoding into three independent 2D rotary subspaces corresponding to the (x, y) , (y, z) , and (z, x) planes. This yields a 6-dimensional geometric embedding structure, organized into three pairs, each pair encoding angular information along one coordinate axis.

The 3D rotary position encoding is applied to the projected query and key vectors. Let $\mathbf{q}^{\text{raw}}, \mathbf{k}^{\text{raw}} \in \mathbb{R}^d$ denote the query and key vectors after linear projection but before positional encoding, associated with 3D coordinate $\mathbf{s}_i = (s_{x_i}, s_{y_i}, s_{z_i})$. We define the rotation operator $R_{\mathbf{s}_x, \mathbf{s}_y, \mathbf{s}_z}$ as:

$$R_{\mathbf{s}_x, \mathbf{s}_y, \mathbf{s}_z} \mathbf{q}^{\text{raw}} = \begin{bmatrix} q_0^{\text{raw}} \\ q_1^{\text{raw}} \\ q_2^{\text{raw}} \\ q_3^{\text{raw}} \\ q_4^{\text{raw}} \\ q_5^{\text{raw}} \end{bmatrix} \odot \begin{bmatrix} \cos(s_x \theta_0) \\ \cos(s_x \theta_0) \\ \cos(s_y \theta_0) \\ \cos(s_y \theta_0) \\ \cos(s_z \theta_0) \\ \cos(s_z \theta_0) \end{bmatrix} + \begin{bmatrix} -q_1^{\text{raw}} \\ q_0^{\text{raw}} \\ -q_3^{\text{raw}} \\ q_2^{\text{raw}} \\ -q_5^{\text{raw}} \\ q_4^{\text{raw}} \end{bmatrix} \odot \begin{bmatrix} \sin(s_x \theta_0) \\ \sin(s_x \theta_0) \\ \sin(s_y \theta_0) \\ \sin(s_y \theta_0) \\ \sin(s_z \theta_0) \\ \sin(s_z \theta_0) \end{bmatrix}, \quad (2)$$

where \odot denotes element-wise multiplication. The final geometry-aware query and key are given by $\mathbf{q} = R_{\mathbf{s}_i} \mathbf{q}^{\text{raw}}$ and $\mathbf{k} = R_{\mathbf{s}_i} \mathbf{k}^{\text{raw}}$, respectively. This formulation ensures that the inner product satisfies the relative positional property: $(R_{\mathbf{s}_1} \mathbf{q}^{\text{raw}})^\top (R_{\mathbf{s}_2} \mathbf{k}^{\text{raw}}) = (\mathbf{q}^{\text{raw}})^\top R_{\mathbf{s}_1 - \mathbf{s}_2} \mathbf{k}^{\text{raw}}$, which encodes 3D spatial relationships directly into attention scores. A detailed derivation is provided in Appendix C.

We implement three RoPE modes that differ in how the input embedding is processed before applying 3D rotary position encoding:

- **Selective:** The input embedding $\mathbf{x}_i \in \mathbb{R}^d$ is split into two halves. The first half (spatial part) is linearly projected and then transformed by 3D RoPE; the second half (feature part) is kept unchanged and directly concatenated.
- **Separate:** The two halves are independently linearly projected, but only the first half is connected through 3D RoPE.
- **Full:** The entire embedding is linearly projected as a vector and fully transformed by 3D RoPE.

Input representation. Each token t_i corresponds to a genomic bin with known 3D coordinates $\mathbf{s}_i = (s_{x_i}, s_{y_i}, s_{z_i}) \in \mathbb{R}^3$. The total number of bins for a chromosome at a given resolution defines the vocabulary size. We map each bin ID $t_i \in \{0, \dots, \text{vocab_size} - 1\}$ to a learnable semantic embedding $\mathbf{E}_{\text{token}}(t_i) \in \mathbb{R}^{d_t}$ via a matrix $\mathbf{W}_{\text{tok}} \in \mathbb{R}^{\text{vocab_size} \times d_t}$.

The initial token representation is formed by concatenating this semantic embedding with the raw spatial coordinates:

$$\mathbf{x}_i = [\mathbf{E}_{\text{token}}(t_i); \mathbf{s}_i] \in \mathbb{R}^d, \quad (3)$$

Let $\mathbf{x}_i^{(s)}$ and $\mathbf{x}_i^{(f)}$ denote the spatial and feature halves of \mathbf{x}_i , each of dimension $d/2$. The 3D rotary position embedding is applied as follows:

$$\text{RoPE-3D}(\mathbf{x}_i) = \begin{cases} [R_{\mathbf{s}_i}(W^{\text{rope}} \mathbf{x}_i^{(s)}); \mathbf{x}_i^{(f)}] & \text{(Selective)} \\ [R_{\mathbf{s}_i}(W^{\text{rope}} \mathbf{x}_i^{(s)}); W^{\text{feat}} \mathbf{x}_i^{(f)}] & \text{(Separate)} \\ R_{\mathbf{s}_i}(W \mathbf{x}_i) & \text{(Full)} \end{cases}, \quad (4)$$

216 where W^{rope} , W^{feat} , W are learnable linear projections, and $R_{\mathbf{s}_i} = R_{s_{x_i}, s_{y_i}, s_{z_i}}$ is the 6D rotary
 217 transformation defined in Equation (2).

218 **(2) Nyström Approximation for Structure Tokenization.** While RoPE-3D in Equation (2) effectively
 219 encodes absolute and relative spatial positions, its axis-wise rotation mechanism is inherently
 220 limited in modeling global pairwise distance relationships, such as long-range structural dependencies
 221 or non-local geometric patterns. To address this, we incorporate a Nyström-based feature en-
 222 coder (Williams & Seeger, 2000; Yang et al., 2012) that explicitly captures low-rank approximations
 223 of the radial basis function (RBF) kernel over 3D coordinates.

224 Formally, let $\mathbf{s}_i \in \mathbb{R}^3$ denote the canonicalized coordinate of token i . We define an RBF kernel
 225 between any two points as:

$$227 \quad \kappa_g(\mathbf{s}_i, \mathbf{s}_j) = \exp\left(-\frac{\|\mathbf{s}_i - \mathbf{s}_j\|^2}{2\sigma_g^2}\right),$$

229 where $\sigma_g > 0$ is the bandwidth of the g -th Gaussian kernel, and we consider a set of G scales
 230 $\{\sigma_g\}_{g=1}^G$. The Nyström method proceeds as follows:

231 **Step 1: Anchor point selection.** We fix a set of m anchor points $\{\mathbf{u}_k\}_{k=1}^m \subset \mathbb{R}^3$, sampled uniformly
 232 from the 3D space. For each scale σ_g , we construct the anchor-anchor Gram matrix $A_g \in \mathbb{R}^{m \times m}$
 233 with entries:

$$234 \quad [A_g]_{k\ell} = \kappa_g(\mathbf{u}_k, \mathbf{u}_\ell) = \exp\left(-\frac{\|\mathbf{u}_k - \mathbf{u}_\ell\|^2}{2\sigma_g^2}\right).$$

235 To ensure numerical stability, we compute the Cholesky decomposition $A_g = O_g O_g^\top$, where $O_g \in$
 236 $\mathbb{R}^{m \times m}$ is a lower-triangular matrix, and precompute $O_g^{-\top}$ for later use.

237 **Step 2: Token–anchor similarity.** For each token coordinate \mathbf{s}_i , we compute its RBF similarities
 238 to all anchors under scale σ_g :

$$239 \quad V_{g,i} = [\kappa_g(\mathbf{s}_i, \mathbf{u}_1), \kappa_g(\mathbf{s}_i, \mathbf{u}_2), \dots, \kappa_g(\mathbf{s}_i, \mathbf{u}_m)] \in \mathbb{R}^m.$$

240 **Step 3: Nyström projection.** We project $V_{g,i}$ using the precomputed inverse factor:

$$241 \quad \tilde{k}_{g,i} = V_{g,i} O_g^{-\top} \in \mathbb{R}^m.$$

242 This yields a low-rank approximation of the full kernel embedding.

243 **Step 4: Multi-scale fusion and compression.** We concatenate the projected features across all
 244 scales:

$$245 \quad \tilde{k}_i = [\tilde{k}_{1,i}; \tilde{k}_{2,i}; \dots; \tilde{k}_{G,i}] \in \mathbb{R}^{Gm},$$

246 and apply a learnable linear projection $f_\theta : \mathbb{R}^{Gm} \rightarrow \mathbb{R}^m$ to obtain the final Nyström structure
 247 embedding:

$$248 \quad \mathbf{E}_{\text{nyström}}(\mathbf{s}_i) = f_\theta(\tilde{k}_i). \quad (5)$$

249 This embedding encodes multi-scale, low-rank geometric information about \mathbf{s}_i and is fused with
 250 token and positional representations in the subsequent transformer layers.

251 3.3 STRUCTURE-AWARE FUSION

252 We fuse the geometry-aware positional encoding and Nyström structure features into a unified Trans-
 253 former input.

254 **Input Representation.** For each token i , we construct the initial embedding by concatenating three
 255 geometric components: (1) the base position embedding \mathbf{x}_i , (2) the normalized canonical coordinate
 256 $\frac{\mathbf{s}_i}{\|\mathbf{s}_i\|}$, which encodes directional information, and (3) the Nyström structure embedding $\mathbf{E}_{\text{nyström}}(\mathbf{s}_i)$.
 257 This yields:

$$258 \quad \mathbf{h}_i^0 = \text{Concat}\left(\mathbf{x}_i, \frac{\mathbf{s}_i}{\|\mathbf{s}_i\|}, \mathbf{E}_{\text{nyström}}(\mathbf{s}_i)\right) \in \mathbb{R}^{d_{\text{in}}}. \quad (6)$$

259 **Position-Augmented Transformer Input.** We then add the geometry-aware positional encoding
 260 RoPE-3D(\mathbf{x}_i) to inject relative spatial context, followed by dropout:

$$261 \quad \mathbf{H}_0 = \text{Dropout}(\mathbf{h}^0 + \text{RoPE-3D}(\mathbf{x}_i)). \quad (7)$$

270 **Transformer Backbone.** The sequence \mathbf{H}_0 is processed by L stacked Transformer layers:
 271

$$272 \quad \mathbf{H}_{l+1} = \text{TransformerLayer}_l(\mathbf{H}_l), \quad l = 0, \dots, L-1. \quad (8)$$

273 The final representation \mathbf{H}_L aggregates multi-scale geometric and structural information, and is used
 274 to predict the 3D coordinates of chromosome bins.
 275

276 **3.4 LEARNING OBJECTIVE**
 277

278 The overall training objective is a hybrid loss that combines a structural-preserving term and a value-
 279 weighted regression term:
 280

$$281 \quad \mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{struct}} + \beta \mathcal{L}_{\text{weighted.mse}}, \quad \beta = 1 - \alpha, \quad (9)$$

282 where $\alpha \in [0, 1]$ balances the two components.
 283

284 **Structural-learning loss.** Let $D = \{D_{i,j}\}$ denote the input pairwise heterogeneity (derived from
 285 the contact-derived distance representation) and let $B = \{\mathbf{b}_i\}$ be the predicted 3D coordinates with
 286 $\mathbf{b}_i \in \mathbb{R}^3$. For each bin i we define neighborhood selection probabilities in input and output spaces:
 287

$$288 \quad p_{j|i} = \frac{\exp(-D_{i,j})}{\sum_{k \neq i} \exp(-D_{i,k})}, \quad q_{j|i} = \frac{\exp(-\|\mathbf{b}_i - \mathbf{b}_j\|^2)}{\sum_{k \neq i} \exp(-\|\mathbf{b}_i - \mathbf{b}_k\|^2)}. \quad (10)$$

289 To align neighborhood structure between input and output we use a bidirectional Kullback–Leibler
 290 (KL) divergence (Gong et al., 2023) objective:
 291

$$292 \quad \mathcal{L}_{\text{struct}} = \lambda \text{KL}(P\|Q) + (1 - \lambda) \text{KL}(Q\|P), \quad (11)$$

293 where $\text{KL}(P\|Q) = \sum_i \sum_{j \neq i} p_{j|i} \log \frac{p_{j|i}}{q_{j|i}}$ and $\lambda \in [0, 1]$ trades off false positives and misses (use
 294 $\lambda = 0.1$ by default). Full derivations are given in Appendix D.1.
 295

296 **Value-weighted MSE.** While structural loss preserves global topology, precise distance prediction
 297 requires an additional term. Hi-C data exhibits higher reliability and biological significance for
 298 smaller distances (high-intensity contacts). We thus introduce a weighted mean squared error (MSE)
 299 (Wang et al., 2024) that assigns adaptive weights based on true value ranks rather than treating all
 300 errors equally. For a batch with n distinct distances, we compute weights w_i for each value (see
 301 Appendix D.2). Then, the weighted MSE is:
 302

$$303 \quad \mathcal{L}_{\text{weighted.mse}} = \sum_{v \in \mathcal{V}} w_v \cdot \frac{1}{N_v} \sum_{(i,j) \in \mathcal{I}_v} (y_{ij} - \hat{y}_{ij})^2, \quad (12)$$

305 where \mathcal{V} indexes distinct true distance values in the batch, \mathcal{I}_v are the pairs with true value v , $N_v =$
 306 $|\mathcal{I}_v|$, y_{ij} is the target distance and \hat{y}_{ij} is the predicted distance.
 307

308 **3.5 STABILITY OF INERTIAL FRAME ALIGNMENT**
 309

310 During our experiments we observed that inputs from physics- or regularization-based reconstruc-
 311 tions (e.g., 3DMax, LorDG) consistently benefited from inertial-frame alignment, whereas contact-
 312 matrix eigendecomposition methods (e.g., Gram) showed little or no gain. To explain this contrast,
 313 recall that the input coordinates are defined as $\mathbf{C}^* = \{(x_i^*, y_i^*, z_i^*)\}_{i=1}^N \in \mathbb{R}^{N \times 3}$. We define its
 314 sample covariance
 315

$$316 \quad \Sigma_{\mathbf{C}^*} = \frac{1}{N} (\mathbf{C}^*)^\top \mathbf{C}^*, \quad (13)$$

317 which admits eigenvalues $\mu_1 \geq \mu_2 \geq \mu_3$ and orthonormal eigenvectors u_1, u_2, u_3 corresponding to
 318 the principal axes. The spectral gap
 319

$$320 \quad \delta(\mathbf{C}^*) = \mu_1 - \mu_2 \quad (14)$$

321 quantifies how well the first principal direction is separated from the remainder.
 322

323 Given two coordinate sets $\mathbf{C}^{*(1)}$ and $\mathbf{C}^{*(2)}$, we measure the angular difference of their leading
 324 inertial axes by
 325

$$326 \quad \theta_{\text{PC1}}(\mathbf{C}^{*(1)}, \mathbf{C}^{*(2)}) = \arccos(|u_1(\mathbf{C}^{*(1)})^\top u_1(\mathbf{C}^{*(2)})|). \quad (15)$$

324 The stability of the principal directions under perturbations is controlled by the Davis–Kahan theorem
 325 (Davis & Kahan, 1970). Let A be a symmetric matrix with leading eigenvector u and spectral
 326 gap δ . If $\tilde{A} = A + \Delta A$ has leading eigenvector \tilde{u} , then
 327

$$328 \quad \sin \angle(u, \tilde{u}) \leq \frac{\|\Delta A\|_2}{\delta}, \quad (16)$$

329
 330

331 where $\|\cdot\|_2$ denotes the spectral norm. Applied to Σ_{C^*} (Equation (13)), Equation (16) shows that
 332 when the spectral gap $\delta(C^*)$ (Equation (14)) is small, even minor perturbations in the data can
 333 rotate the leading axis significantly. Conversely, a large gap yields stable inertial axes and reliable
 334 alignment. This analysis and explanation can be found in Appendix E.
 335
 336

337 4 EXPERIMENTS

338

339 **Datasets.** We evaluated our method on two single-cell Hi-C datasets: human frontal cortex cell
 340 (Wang & Cheng, 2024) and B-Lymphocyte cell (Oluwadare et al., 2020). Both datasets followed the
 341 chromosome partitioning scheme of Wang & Cheng (2025), with training sets (frontal cortex: chr
 342 1,3,5,7,8,9,11,13,15,16,17,19,21,22; B-Lymphocyte: adds chr 23), validation sets (chr 2,6,10,12 for
 343 both), and test sets (chr 4,14,18,20 for both).

344 **Baselines and Implementation.** We compare our method with both classical numerical and deep
 345 learning baselines for 3D chromosome structure reconstruction. The classical numerical methods
 346 include 3DMax (Oluwadare et al., 2018) and LorDG (Trieu & Cheng, 2017), while the deep learning
 347 baselines include HiC-GNN (Hovenga et al., 2023) and HiCEGNN (Wang & Cheng, 2025). All
 348 baselines were run with their default configurations as provided in the respective source codes.
 349

350 **Metrics.** We evaluate 3D chromosome reconstruction using two metrics: (1) **Distance Spearman**
 351 **correlation coefficient (dSCC)** (Oluwadare et al., 2018) measures rank correlation between pre-
 352 dicted and ideal distances (range $[-1,1]$); higher values indicate better structural quality and scale
 353 invariance. (2) **Distance root mean square error (dRMSE)** (Varoquaux et al., 2014) quantifies
 354 absolute distance errors; lower values denote higher accuracy and similarity to ideal distance map.
 355

356 4.1 MAIN RESULTS

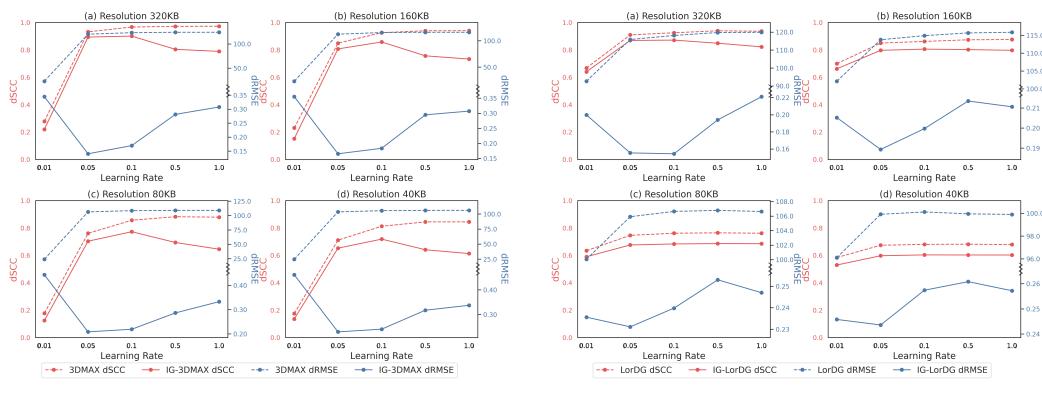
357

358 **Reconstruction performance on Frontal cortex cell dataset.** Table 1 compares six methods on
 359 single-cell Hi-C data. Our two variants, IG-3DMAX and IG-LorDG, consistently outperform all
 360 baselines in both dSCC and dRMSE across all resolutions. For example, at 320 kB, IG-3DMAX
 361 achieves a dSCC of 0.9006, significantly higher than HiCEGNN (0.5804) and 3DMAX (0.2780),
 362 while reducing dRMSE to 0.1697 from 0.2744. Similar improvements are observed at other reso-
 363 lutions, with dSCC gains often exceeding 50% and dRMSE reductions of 30–40%. These results
 364 demonstrate that combining inertial-frame canonicalization with our Transformer yields state-of-
 365 the-art accuracy and stronger resolution-agnostic performance compared to equivariant and num-
 366 erical baselines. Traditional numerical methods (3DMAX, LorDG) perform substantially worse, with
 367 dRMSE values orders of magnitude higher, due to the lack of direct distance supervision. Gram-
 368 matrix inputs also yield poor performance in our pipeline; we analyze this instability in Appendix E.
 369

370 **Reconstruction performance on B-Lymphocyte cell dataset.** Table 5 compares six methods on
 371 single-cell Hi-C data. IG-3DMAX achieves the best dSCC and lowest dRMSE across all four reso-
 372 lutions, demonstrating robust superiority. For example, at 1MB it attains a dSCC of 0.9209 and
 373 dRMSE of 0.0822, outperforming all baselines. Similar advantages are observed at finer resolutions
 374 (500KB, 250KB, 100KB), with dSCC consistently above 0.87 and dRMSE below 0.08. IG-LorDG
 375 shows competitive dSCC at 500KB (0.8367) and 250KB (0.8440), but weaker performance at other
 376 resolutions. This variation stems from LorDG’s inherent reconstruction limitations on this dataset
 377 (e.g., high dRMSE 107.7091 at 1MB), which constrain IG-LorDG’s input quality. In contrast, IG-
 378 3DMAX’s stability highlights our framework’s resilience to input variations, delivering state-of-the-
 379 art accuracy regardless of baseline method performance (see Appendix F).

378
379 Table 1: Performance comparison of six methods on 3D chromosome structure reconstruction from
380 single-cell Hi-C data (Frontal cortex cell test set). Metrics report distance-based Spearman correla-
381 tion (dSCC \uparrow) and root mean square error (dRMSE \downarrow) at four resolutions. Best results in bold.
382

Method	320KB		160KB		80KB		40KB	
	dSCC \uparrow	dRMSE \downarrow						
3DMAX	0.2780	23.1587	0.2302	23.3439	0.1774	23.9174	0.1754	24.6538
LorDG	0.6681	92.6582	0.6997	102.1392	0.6342	100.0507	0.5841	96.1048
HiC-GNN	0.2432	0.8366	0.2077	0.9083	0.1370	0.9352	0.0915	0.9456
HiCEGNN	0.5804	0.2744	0.5351	0.3550	0.3288	0.4158	0.2506	0.4317
IG-3DMAX	0.9006	0.1697	0.8577	0.1835	0.7727	0.2192	0.7187	0.2410
IG-LorDG	0.8713	0.1544	0.8056	0.1997	0.6835	0.2398	0.6036	0.2574



403
404 Figure 3: Learning rate ablation: performance at four resolutions for five learning rates. Red: dSCC
405 (left axis). Blue: dRMSE (right axis).
406

4.2 ABLATION STUDIES

410 **Learning rate robustness.** We evaluate our method’s robustness to input quality by varying the
411 learning rates $\{1, 0.5, 0.1, 0.05, 0.01\}$ of 3DMAX and LorDG on Frontal cortex data. Input coordi-
412 nates generated under each setting are evaluated at four resolutions using dSCC (\uparrow) and dRMSE (\downarrow).
413 Figure 3 shows: (1) 3DMAX vs. IG-3DMAX, (2) LorDG vs. IG-LorDG. Our IG variants maintain
414 consistently high dSCC while reducing dRMSE to 10^{-1} scale across all learning rates. Baselines show
415 significant dRMSE fluctuations (up to 10^2 scale) and unstable performance. This demonstrates
416 InertialGenome’s superior stability and accuracy regardless of input learning rate and resolution.

417 **Loss Components** We analyze the contribution of the structural stability loss ($\mathcal{L}_{\text{struct}}$) versus the
418 coordinate regression loss ($\mathcal{L}_{\text{weighted_mse}}$) by varying their weighting ratio α/β in the total objective
419 (see Appendix G.1).

420 **Component ablation.** We assess the impact of key design choices by removing inertial-frame
421 alignment, RoPE-3D, or Nyström encoding from our full model (see Appendix G.2).

4.3 CASE STUDIES

425 **A/B compartment validation.** We validate the biological plausibility of our reconstructed struc-
426 tures using A/B compartment analysis (see Appendix H.1). Appendix Figure 6 compares dis-
427 tance distributions for IG-3DMAX and HICEGNN. IG-3DMAX shows significantly shorter intra-
428 compartment (A–A, B–B) than inter-compartment (A–B) distances ($p_A = 0.0001$, $p_B = 0.0038$),
429 confirming expected compartmental organization. In contrast, HICEGNN shows no significant
430 A–A/A–B separation ($p_A = 0.4360$) and weak B–B separation ($p_B = 0.0000$), indicating poor
431 compartmentalization. These results demonstrate that IG-3DMAX better captures spatial compart-
432 ment segregation than HICEGNN.

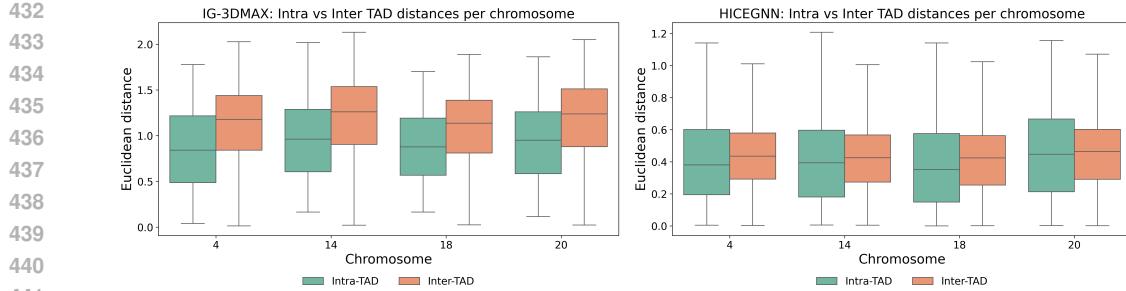


Figure 4: Intra- vs. inter-TAD Euclidean distances for IG-3DMAX (left) and HiCEGNN (right). Lower intra-TAD distances indicate stronger spatial clustering.

Table 2: Cross-resolution transfer results. **Bold**: improvement over same-resolution original model; underline: HiCEGNN transfer improvement.

Method	320kb → 160kb		320kb → 80kb		320kb → 40kb	
	dSCC ↑	dRMSE ↓	dSCC ↑	dRMSE ↓	dSCC ↑	dRMSE ↓
HICEGNN-Transfer	<u>0.5815</u>	0.2997	<u>0.4080</u>	<u>0.4108</u>	0.2343	<u>0.3840</u>
HICEGNN-Original	0.5351	0.3550	0.3288	0.4158	0.2506	0.4317
IG-3DMAX-Full	0.7455	0.3159	0.6766	0.3464	0.6528	0.3588
IG-3DMAX-Selective	0.7399	0.3089	0.6698	0.3447	0.6467	0.3567
IG-3DMAX-Separate	0.7434	0.3057	0.6714	0.3392	0.6480	0.3580
IG-3DMAX-Original	0.7332	0.3071	0.6451	0.3328	0.6132	0.3374

TAD domain consistency validation. We assess whether bin within the same TAD are spatially closer than those across TADs (see Appendix H.2). Figure 4 shows intra- vs. inter-TAD distances for IG-3DMAX and HiCEGNN. IG-3DMAX yields consistently shorter intra-TAD distances across chromosomes 4, 14, 18, and 20, with intra/inter ratios of 0.76–0.80 and highly significant p -values (Appendix Table 8). HiCEGNN shows higher ratios (0.91–0.99) and weaker significance (e.g., $p = 0.159$ for chr20), indicating poor domain compaction. These results confirm that IG-3DMAX better captures TAD-level spatial organization than HiCEGNN.

Cross-resolution transfer learning. We evaluate transfer from 320kb to finer resolutions (160kb, 80kb, 40kb). Table 2 shows IG-3DMAX consistently outperforms both its original model and HICEGNN variants in dSCC and dRMSE. At higher resolutions (80kb, 40kb), IG-3DMAX maintains stable dRMSE (10^{-2}) and improves dSCC by 5%, while HICEGNN degrades (dSCC drops at 40kb). Our gains stem from RoPE and inertial-frame alignment, which preserve spatial relations across scales. HICEGNN lacks such geometric adaptation, leading to unstable performance. IG-3DMAX demonstrates robust, resolution-agnostic reconstruction, with advantages magnified at finer resolutions.

5 CONCLUSION

We presented InertialGenome, a Transformer-based framework for robust, resolution-agnostic 3D chromosome reconstruction. It uses inertial-frame canonicalization for pose invariance and a geometry-aware Transformer with Nyström positional encoding to capture long-range interactions efficiently. Experiments on two single-cell datasets across four resolutions show that InertialGenome consistently outperforms classical (3DMAX, LorDG) and deep learning (HiC-GNN, HiCEGNN) methods in both structural accuracy and functional validation. Our cross-resolution strategy further boosts high-resolution performance by up to 5%, demonstrating strong generalization and biological plausibility.

By decoupling physical constraints from architecture, InertialGenome offers a flexible alternative to $SE(3)$ -equivariant models, paving the way for scalable 3D genome modeling. Future work will integrate multi-modal genomic data to enhance reconstruction robustness across diverse experimental conditions.

486 ETHICS STATEMENT
487488 **Answer:** Our study complies strictly with the ICLR Code of Ethics. The research involves no
489 human participants or sensitive personal data. All datasets and code employed adhere to their re-
490 spective licensing agreements. This work presents foundational research that does not raise concerns
491 related to fairness, privacy, security, or potential misuse. We affirm that all ethical aspects have been
492 comprehensively considered and addressed.493
494 REPRODUCIBILITY STATEMENT
495496 **Answer:** We are dedicated to facilitating the reproducibility of our work. Full details required to
497 replicate our key findings—including data access instructions, experimental parameters, model ar-
498 chitectures, evaluation metrics, and model checkpoints—will be made available on GitHub upon
499 paper acceptance. Users can utilize our detailed documentation and code scripts to faithfully repro-
500 duce the results, ensuring research transparency and methodological rigor.501
502 REFERENCES
503504 Ahmed Abbas, Xuan He, Jing Niu, Bin Zhou, Guangxiang Zhu, Tszshan Ma, Jiangpeikun Song,
505 Juntao Gao, Michael Q Zhang, and Jianyang Zeng. Integrating hi-c and fish data for modeling of
506 the 3d organization of chromosomes. *Nature communications*, 10(1):2049, 2019.507 Badri Adhikari, Tuan Trieu, and Jianlin Cheng. Chromosome3d: reconstructing three-dimensional
508 chromosomal structures from hi-c interaction frequency data using distance geometry simulated
509 annealing. *BMC genomics*, 17(1):886, 2016.510 Mariano Barbieri, Mita Chotalia, James Fraser, Liron-Mark Lavitas, Josée Dostie, Ana Pombo, and
511 Mario Nicodemi. Complexity of chromatin folding is captured by the strings and binders switch
512 model. *Proceedings of the National Academy of Sciences*, 109(40):16173–16178, 2012.513 Robert A Beagrie, Antonio Scialdone, Markus Schueler, Dorothee CA Kraemer, Mita Chotalia,
514 Sheila Q Xie, Mariano Barbieri, Inês de Santiago, Liron-Mark Lavitas, Miguel R Branco, et al.
515 Complex multi-enhancer contacts captured by genome architecture mapping. *Nature*, 543(7646):
516 519–524, 2017.517 Jon-Matthew Belton, Rachel Patton McCord, Johan Harmen Gibcus, Natalia Naumova, Ye Zhan,
518 and Job Dekker. Hi-c: a comprehensive technique to capture the conformation of genomes.
519 *Methods*, 58(3):268–276, 2012.520 Claudia Caudai, Emanuele Salerno, Monica Zoppe, Ivan Merelli, and Anna Tonazzini. Chromstruct
521 4: a python code to estimate the chromatin structure from hi-c data. *IEEE/ACM transactions on*
522 *computational biology and bioinformatics*, 16(6):1867–1878, 2018.523 Nucleome Consortium et al. An integrated view of the structure and function of the human 4d
524 nucleome. *bioRxiv*, 2024.525 Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii.
526 *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.527 Job Dekker, Karsten Rippe, Martijn Dekker, and Nancy Kleckner. Capturing chromosome confor-
528 mation. *science*, 295(5558):1306–1311, 2002.529 Jesse R Dixon, Siddarth Selvaraj, Feng Yue, Audrey Kim, Yan Li, Yin Shen, Ming Hu, Jun S Liu,
530 and Bing Ren. Topological domains in mammalian genomes identified by analysis of chromatin
531 interactions. *Nature*, 485(7398):376–380, 2012.532 Josée Dostie, Todd A Richmond, Ramy A Arnaout, Rebecca R Selzer, William L Lee, Tracey A
533 Honan, Eric D Rubio, Anton Krumm, Justin Lamb, Chad Nusbaum, et al. Chromosome confor-
534 mation capture carbon copy (5c): a massively parallel solution for mapping interactions between
535 genomic elements. *Genome research*, 16(10):1299–1309, 2006.

540 Darya Filippova, Rob Patro, Geet Duggal, and Carl Kingsford. Identification of alternative topolog-
 541 ical domains in chromatin. *Algorithms for Molecular Biology*, 9(1):14, 2014.

542

543 Melissa J Fullwood, Mei Hui Liu, You Fu Pan, Jun Liu, Han Xu, Yusoff Bin Mohamed, Yuriy L
 544 Orlov, Stoyan Velkov, Andrea Ho, Poh Huay Mei, et al. An oestrogen-receptor- α -bound human
 545 chromatin interactome. *Nature*, 462(7269):58–64, 2009.

546

547 Haiyan Gong, Fuqiang Ma, Xiaotong Zhang, Yi Yang, Minghong Li, Zhengyuan Chen, Sichen
 548 Zhang, and Yang Chen. A 3-d chromosome structure reconstruction method with high resolu-
 549 tion hi-c data using nonlinear dimensionality reduction and divide-and-conquer strategy. *IEEE
 550 Transactions on NanoBioscience*, 22(4):716–727, 2023.

551

552 Van Hovenga, Jugal Kalita, and Oluwatosin Oluwadare. Hic-gnn: A generalizable model for
 553 3d chromosome reconstruction using graph convolutional neural networks. *Computational and
 554 Structural Biotechnology Journal*, 21:812–836, 2023.

555

556 Ming Hu, Ke Deng, Zhaohui Qin, Jesse Dixon, Siddarth Selvaraj, Jennifer Fang, Bing Ren, and
 557 Jun S Liu. Bayesian inference of spatial organizations of chromosomes. *PLoS computational
 558 biology*, 9(1):e1002893, 2013.

559

560 Nan Hua, Harianto Tjong, Hanjun Shin, Ke Gong, Xianghong Jasmine Zhou, and Frank Alber.
 561 Producing genome structure populations with the dynamic and automated pgs software. *Nature
 562 protocols*, 13(5):915–926, 2018.

563

564 Fang-Zhen Li, Zhi-E Liu, Xiu-Yuan Li, Li-Mei Bu, Hong-Xia Bu, Hui Liu, and Cai-Ming Zhang.
 565 Chromatin 3d structure reconstruction with consideration of adjacency relationship among ge-
 566 nomic loci. *BMC bioinformatics*, 21(1):272, 2020.

567

568 Erez Lieberman-Aiden, Nynke L Van Berkum, Louise Williams, Maxim Imakaev, Tobias Ragoczy,
 569 Agnes Telling, Ido Amit, Bryan R Lajoie, Peter J Sabo, Michael O Dorschner, et al. Compre-
 570 hensive mapping of long-range interactions reveals folding principles of the human genome. *science*,
 571 326(5950):289–293, 2009.

572

573 Dario Meluzzi and Gaurav Arya. Recovering ensembles of chromatin conformations from contact
 574 probabilities. *Nucleic acids research*, 41(1):63–75, 2013.

575

576 Oluwatosin Oluwadare, Yuxiang Zhang, and Jianlin Cheng. A maximum likelihood algorithm for
 577 reconstructing 3d structures of human chromosomes from chromosomal contact data. *BMC ge-
 578 nomics*, 19(1):161, 2018.

579

580 Oluwatosin Oluwadare, Max Highsmith, Douglass Turner, Erez Lieberman Aiden, and Jianlin
 581 Cheng. Gsdb: a database of 3d chromosome and genome structures reconstructed from hi-c
 582 data. *BMC molecular and cell biology*, 21(1):60, 2020.

583

584 Ana Pombo and Mario Nicodemi. Physical mechanisms behind the large scale features of chromatin
 585 organization. *Transcription*, 5(2):e28447, 2014.

586

587 Sofia A Quinodoz, Prashant Bhat, Peter Chovanec, Joanna W Jachowicz, Noah Ollikainen, Elizabeth
 588 Detmar, Elizabeth Soehalim, and Mitchell Guttman. Sprite: a genome-wide method for mapping
 589 higher-order 3d interactions in the nucleus using combinatorial split-and-pool barcoding. *Nature
 590 protocols*, 17(1):36–75, 2022.

591

592 Suhas SP Rao, Miriam H Huntley, Neva C Durand, Elena K Stamenova, Ivan D Bochkov, James T
 593 Robinson, Adrian L Sanborn, Ido Machol, Arina D Omer, Eric S Lander, et al. A 3d map of
 594 the human genome at kilobase resolution reveals principles of chromatin looping. *Cell*, 159(7):
 595 1665–1680, 2014.

596

597 Lila Rieber and Shaun Mahony. minimds: 3d structural inference from high-resolution hi-c data.
 598 *Bioinformatics*, 33(14):i261–i266, 2017.

599

600 Mathieu Rousseau, James Fraser, Maria A Ferraiuolo, Josée Dostie, and Mathieu Blanchette. Three-
 601 dimensional modeling of chromatin structure from interaction frequency data using markov chain
 602 monte carlo sampling. *BMC bioinformatics*, 12(1):414, 2011.

594 François Serra, Davide Baù, Mike Goodstadt, David Castillo, Guillaume J Filion, and Marc A Marti-
 595 Renom. Automatic analysis and 3d-modelling of hi-c data using tadbit reveals structural features
 596 of the fly chromatin colors. *PLoS computational biology*, 13(7):e1005665, 2017.

597

598 Hanjun Shin, Yi Shi, Chao Dai, Harianto Tjong, Ke Gong, Frank Alber, and Xianghong Jasmine
 599 Zhou. Topdom: an efficient and deterministic method for identifying topological domains in
 600 genomes. *Nucleic acids research*, 44(7):e70–e70, 2016.

601 Marieke Simonis, Petra Klous, Erik Splinter, Yuri Moshkin, Rob Willemsen, Elzo De Wit, Bas
 602 Van Steensel, and Wouter De Laat. Nuclear organization of active and inactive chromatin domains
 603 uncovered by chromosome conformation capture-on-chip (4c). *Nature genetics*, 38(11):1348–
 604 1354, 2006.

605 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: En-
 606 hanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

607

608 Tuan Trieu and Jianlin Cheng. Large-scale reconstruction of 3d structures of human chromosomes
 609 from chromosomal contact data. *Nucleic acids research*, 42(7):e52–e52, 2014.

610 Tuan Trieu and Jianlin Cheng. 3d genome structure modeling by lorentzian objective function.
 611 *Nucleic acids research*, 45(3):1049–1058, 2017.

612

613 Nelle Varoquaux, Ferhat Ay, William Stafford Noble, and Jean-Philippe Vert. A statistical approach
 614 for inferring the 3d structure of the genome. *Bioinformatics*, 30(12):i26–i33, 2014.

615

616 Xiao Wang, Yuanyuan Zhang, Suhita Ray, Anupama Jha, Tangqi Fang, Shengqi Hang, Sergei Doula-
 617 tov, William Stafford Noble, and Sheng Wang. A generalizable hi-c foundation model for chro-
 618 matin architecture, single-cell and multi-omics analysis across species. *bioRxiv*, 2024.

619 Yanli Wang and Jianlin Cheng. Hicdiff: single-cell hi-c data denoising with diffusion models.
 620 *Briefings in Bioinformatics*, 25(4), 2024.

621

622 Yanli Wang and Jianlin Cheng. Reconstructing 3d chromosome structures from single-cell hi-c data
 623 with so (3)-equivariant graph neural networks. *NAR Genomics and Bioinformatics*, 7(1):lqaf027,
 624 2025.

625 Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines.
 626 *Advances in neural information processing systems*, 13, 2000.

627

628 Joachim Wolff, Leily Rabbani, Ralf Gilsbach, Gautier Richard, Thomas Manke, Rolf Backofen, and
 629 Björn A Grüning. Galaxy hicexplorer 3: a web server for reproducible hi-c, capture hi-c and
 630 single-cell hi-c data analysis, quality control and visualization. *Nucleic acids research*, 48(W1):
 631 W177–W184, 2020.

632 Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs
 633 random fourier features: A theoretical and empirical comparison. *Advances in neural information
 634 processing systems*, 25, 2012.

635

636 ZhiZhuo Zhang, Guoliang Li, Kim-Chuan Toh, and Wing-Kin Sung. Inference of spatial orga-
 637 nizations of chromosomes using semi-definite embedding approach and hi-c data. In *Annual
 638 international conference on research in computational molecular biology*, pp. 317–332. Springer,
 639 2013.

640

641

642

643

644

645

646

647

648 A THE USE OF LARGE LANGUAGE MODELS
649650 **Answer:** In the preparation of this work, the authors used large language models (LLMs) for two
651 specific purposes: (1) to refine and polish English language expression, and (2) to assist in the
652 formulation and typesetting of mathematical equations in LaTeX. All scientific content, research
653 design, analysis, and conclusions remain solely the responsibility of the authors.
654655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 **B RELATED WORK**
703704 **B.1 DISTANCE AND CONTACT CONSTRAINED OPTIMIZATION METHODS**
705

706 Distance-based methods rely on the assumption that genomic loci closer in the one-dimensional
707 sequence are also spatially proximal in three dimensions. The central idea is to exploit the inverse
708 relationship between Hi-C contact frequency and physical distance, thereby transforming the contact
709 matrix into a distance matrix and reformulating the reconstruction task as inferring 3D coordinates
710 from pairwise distances. A variety of algorithms adopt this paradigm, including miniMDS (Rieber
711 & Mahony, 2017), 3DMax (Oluwadare et al., 2018), ChromSDE (Zhang et al., 2013), Chromo-
712 some3D (Adhikari et al., 2016), ShNeigh1 (Li et al., 2020), LorDG (Trieu & Cheng, 2017), and
713 others. These methods aim to faithfully recover the underlying geometric distances. For example,
714 3DMax (Oluwadare et al., 2018) assumes that Hi-C data follow a Gaussian distribution and that
715 contact counts are independent given the structure, defining a log-likelihood objective to identify
716 the most probable conformation. LorDG (Trieu & Cheng, 2017), in contrast, employs a nonlinear
717 Lorentzian function to enforce structural consistency while improving robustness to noisy distance
718 constraints. In contrast, contact-based methods (Meluzzi & Arya, 2013; Trieu & Cheng, 2014; Ab-
719bas et al., 2019) directly translate interaction frequencies (IFs) into spatial constraints without first
720 converting them into distances. For example, GEM (Abbas et al., 2019) enforces distance thresholds
721 between chromosome segments and incorporates biophysical feasibility to reconstruct 3D structures.

722 **B.2 PROBABILITY BASED METHODS**
723

724 Probabilistic methods typically assume that the global 3D structure underlies the observed contact
725 map and formulate chromosome reconstruction as a Bayesian inference or maximum likelihood
726 problem. In this framework, the contact matrix is modeled as data generated from specific prob-
727 ability distributions, with representative approaches including BACH (Hu et al., 2013), MCMC5C
728 (Rousseau et al., 2011), PASTIS (Varoquaux et al., 2014), PGS (Hua et al., 2018), and CHROM-
729 STRUCT 4 (Caudai et al., 2018). For instance, MCMC5C (Rousseau et al., 2011) employs a Gaus-
730 sian prior and uses Markov Chain Monte Carlo (MCMC) sampling to infer spatial coordinates from
731 the posterior distribution of interaction frequencies. BACH (Hu et al., 2013) assumes a Poisson
732 distribution for contact counts and applies MCMC to sample chromosome conformations. PASTIS
733 (Varoquaux et al., 2014), in contrast, optimizes spatial coordinates by maximizing the likelihood
734 under the assumption that contact frequencies follow a Poisson distribution conditioned on 3D po-
735 sitions.

736 **B.3 DEEP LEARNING BASED METHODS**
737

738 Unlike traditional distance and probability-based approaches, deep learning offers a fundamentally
739 different paradigm for reconstructing 3D chromosome structures from Hi-C data. HiC-GNN (Hov-
740 enga et al., 2023) pioneered this line of work by applying graph convolutional networks to capture
741 neighborhood structures in Hi-C interaction graphs, directly inferring 3D conformations from large-
742 scale data. Recently, HiCEGNN (Wang & Cheng, 2025) extended this idea by employing an $SO(3)$ -
743 equivariant graph neural network (EGNN) to account for translational and rotational symmetries,
744 enabling accurate prediction of the 3D coordinates of genomic loci. These graph-based deep learn-
745 ing approaches demonstrate not only the feasibility but also the advantages of deep learning for 3D
746 genome modeling: they are faster and more straightforward than traditional optimization methods,
747 and they can leverage large Hi-C datasets to capture common structural patterns often missed by
748 conventional approaches. Building on this line of work, we propose a new deep learning frame-
749 work. In contrast to existing graph-based methods, our approach demonstrates that a Transformer
750 architecture can provide superior modeling capacity for 3D chromosome structure prediction.

751 **C DERIVATION OF 3D GEOMETRIC POSITION EMBEDDING**
752

753 We aim to extend the rotary position embedding (RoPE) mechanism to 3D Euclidean space for
754 chromosome structure modeling. The goal is to design a geometric position embedding such that
755 the relative rotation between two positions satisfies:

$$R_{x_1, y_1, z_1}^\top R_{x_2, y_2, z_2} = R_{x_1 - x_2, y_1 - y_2, z_1 - z_2}. \quad (17)$$

To achieve this, we first recall that any imaginary number $z = a + bi$ can be written in polar form as:

$$z = r \cos \theta + ir \sin \theta = re^{i\theta}, \quad (18)$$

where $r = |z|$ and $\theta = \arg(z)$. This can further be expressed in matrix exponential form using the identity:

$$r \exp(\theta J) \equiv r \exp(\theta) \cdot J, \quad (19)$$

with $J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$, which satisfies $J^2 = -I$, $J^3 = -J$, $J^4 = I$. Thus,

$$r \exp(\theta J) = r \sum_{n=0}^{\infty} \frac{(\theta J)^n}{n!} = r \left(I + \theta J - \frac{\theta^2}{2!} I - \frac{\theta^3}{3!} J + \dots \right). \quad (20)$$

Grouping terms by even and odd powers:

$$r \exp(\theta J) = r \left(\sum_{n=0}^{\infty} \frac{(-1)^n \theta^{2n}}{(2n)!} I \right) + r \left(\sum_{n=0}^{\infty} \frac{(-1)^n \theta^{2n+1}}{(2n+1)!} J \right) \quad (21)$$

$$= r \cos \theta \cdot I + r \sin \theta \cdot J \quad (22)$$

$$= r \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (23)$$

Hence, the 2D rotation matrix can be written as $R_m = \exp(mJ)$, and it follows that:

$$R_m^\top R_n = \exp(-mJ) \exp(nJ) = \exp((n-m)J) = R_{n-m}. \quad (24)$$

Now, we generalize this to 3D. We define a 3D rotation operator $R_{x,y,z}$ as the exponential of a skew-symmetric matrix:

$$R_{x,y,z} = \exp \left(\theta \begin{bmatrix} 0 & -x & 0 & 0 & 0 & 0 \\ x & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -y & 0 & 0 \\ 0 & 0 & y & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -z \\ 0 & 0 & 0 & 0 & z & 0 \end{bmatrix} \right), \quad (25)$$

which decomposes into three independent 2D rotations along the (x, y) , (y, z) , and (z, x) planes.

Expanding this matrix exponential yields:

$$R_{x,y,z} = \begin{bmatrix} \cos x\theta & -\sin x\theta & 0 & 0 & 0 & 0 \\ \sin x\theta & \cos x\theta & 0 & 0 & 0 & 0 \\ 0 & 0 & \cos y\theta & -\sin y\theta & 0 & 0 \\ 0 & 0 & \sin y\theta & \cos y\theta & 0 & 0 \\ 0 & 0 & 0 & 0 & \cos z\theta & -\sin z\theta \\ 0 & 0 & 0 & 0 & \sin z\theta & \cos z\theta \end{bmatrix}. \quad (26)$$

When applied to a 6-dimensional token embedding $\mathbf{q} = [q_0, q_1, q_2, q_3, q_4, q_5]^\top$, we obtain:

$$R_{x,y,z} \mathbf{q} = \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \\ q_4 \\ q_5 \end{bmatrix} \odot \begin{bmatrix} \cos x\theta_0 \\ \cos x\theta_0 \\ \cos y\theta_0 \\ \cos y\theta_0 \\ \cos z\theta_0 \\ \cos z\theta_0 \end{bmatrix} + \begin{bmatrix} -q_1 \\ q_0 \\ -q_3 \\ q_2 \\ -q_5 \\ q_4 \end{bmatrix} \odot \begin{bmatrix} \sin x\theta_0 \\ \sin x\theta_0 \\ \sin y\theta_0 \\ \sin y\theta_0 \\ \sin z\theta_0 \\ \sin z\theta_0 \end{bmatrix}, \quad (27)$$

where θ_0 is the base frequency (e.g., $10000^{-2/6}$), and \odot denotes element-wise multiplication.

This construction ensures that the inner product satisfies:

$$(R_{x_1, y_1, z_1} \mathbf{q})^\top (R_{x_2, y_2, z_2} \mathbf{k}) = \mathbf{q}^\top R_{x_1 - x_2, y_1 - y_2, z_1 - z_2} \mathbf{k}, \quad (28)$$

as required. Therefore, our 6D geometric embedding naturally captures relative spatial relationships through rotational invariance.

810 **D DETAILED LOSS FUNCTIONS**
811812 **D.1 STRUCTURAL-LEARNING LOSS**
813814 Let the spatial distance representation be
815

816
$$D = \{D_1, D_2, \dots, D_n\}, \quad D_i = \{D_{i,1}, D_{i,2}, \dots, D_{i,n}\}, \quad (29)$$

817 and the 3D coordinate representation be
818

819
$$S = \{S_1, S_2, \dots, S_n\}, \quad S_i = (x_i, y_i, z_i). \quad (30)$$

820 **Neighborhood Probability Distributions.** Based on heterogeneity between genomic bins, the
821 probabilities are:
822

823
$$p_{j|i} = \frac{\exp(-D_{ij})}{\sum_{k \neq i} \exp(-D_{ik})}, \quad q_{j|i} = \frac{\exp(-\|\mathbf{s}_i - \mathbf{s}_j\|^2)}{\sum_{k \neq i} \exp(-\|\mathbf{s}_i - \mathbf{s}_k\|^2)}, \quad (31)$$

825 where $\|\mathbf{s}_i - \mathbf{s}_j\|^2$ is the squared Euclidean distance between bins i and j .
826827 **Bidirectional KL Divergence.** Let $P = \{P_1, \dots, P_n\}$ be the distance distribution and $Q =$
828 $\{Q_1, \dots, Q_n\}$ the 3D-space distribution. We define
829

830
$$KL(P\|Q) = \sum_i \sum_{j \neq i} p_{j|i} \log \frac{p_{j|i}}{q_{j|i}}, \quad (32)$$

832
$$KL(Q\|P) = \sum_i \sum_{j \neq i} q_{j|i} \log \frac{q_{j|i}}{p_{j|i}}. \quad (33)$$

834 Balancing false positives and misses via parameter λ , the final structural-learning loss is:
835

836
$$\begin{aligned} \mathcal{L}_{\text{struct}} &= \lambda KL(P\|Q) + (1 - \lambda) KL(Q\|P) \\ &= \lambda \sum_i \sum_{j \neq i} p_{j|i} \log \frac{p_{j|i}}{q_{j|i}} + (1 - \lambda) \sum_i \sum_{j \neq i} q_{j|i} \log \frac{q_{j|i}}{p_{j|i}}. \end{aligned} \quad (34)$$

840 **D.2 VALUE-WEIGHTED MSE LOSS**
841842 To emphasize smaller distances (high-intensity contacts), we weight the MSE by ranks. For a batch
843 with n distinct true distances, the weight for each value is computed as:
844

845
$$w_i = \frac{\text{rank}(i)}{n(n+1)/2}, \quad (35)$$

847 where $\text{rank}(i)$ is the ascending rank order of value i (smallest value has rank 1 and receives the
848 highest weight).
849850 The final weighted MSE loss is:
851

852
$$\mathcal{L}_{\text{weighted-mse}} = \sum_{i \in \text{values}} w_i \cdot \frac{\sum (y_i - p_i)^2}{N_i}, \quad (36)$$

853 where N_i is the frequency of true value i in the current batch.
854855 By combining Equation (34) and Equation (36), our model simultaneously learns accurate local
856 distances and a consistent global structure, enabling more accurate and robust 3D chromosome
857 reconstruction.
858859 **E DERIVATIONS FOR INERTIAL ALIGNMENT STABILITY**
860861 We summarize here the derivation and the empirical checks that complement the main text.
862863 **Notation.** $C \in \mathbb{R}^{N \times N}$ denotes the contact matrix, $D = f(C)$ its corresponding distance matrix ob-
tained through the monotone transform $f(\cdot)$, $H = I - \frac{1}{N} \mathbf{1} \mathbf{1}^\top$ the centering matrix, $B = -\frac{1}{2} H D^2 H$

864 the Gram matrix, and $X \in \mathbb{R}^{N \times 3}$ the centered 3D coordinates whose sample covariance is
 865 $\Sigma_X = \frac{1}{N} X^\top X$.

866 **First-order perturbation.** Let ΔC be a small perturbation of C . By Taylor expansion of f we have
 867 to first order

$$868 \quad \Delta D \approx f'(C) \circ \Delta C.$$

870 Consequently, the Gram perturbation is

$$871 \quad \Delta B \approx -H(D \circ \Delta D)H.$$

873 Taking operator norms and using $\|H\|_2 = 1$ yields the bound

$$874 \quad \|\Delta B\|_2 \lesssim \|D\|_\infty \|f'(C)\|_\infty \|\Delta C\|_F.$$

876 Thus the size of ΔB scales linearly with the contact perturbation norm and with the current dis-
 877 tances.

878 **Effect on principal components.** Let u_1 be the leading eigenvector of B and \tilde{u}_1 that of $B + \Delta B$.
 879 By the Davis–Kahan theorem (Equation (16) in the main text) we have

$$880 \quad \sin \angle(u_1, \tilde{u}_1) \leq \frac{\|\Delta B\|_2}{\delta_B},$$

883 where δ_B is the spectral gap of B (difference between its first and second eigenvalues). Hence, even
 884 very small perturbations of C can produce large rotations of the leading direction whenever δ_B is
 885 small; conversely, a large gap yields robust orientation.

886 **Empirical validation (Chromosome 3 at 320 kb).** We found through numerical calculations that
 887 many of the chromosome 3D coordinates used as input in grams are coplanar or almost coplanar,
 888 which directly indicates the problem of using grams as input. Of course, we still validated the
 889 difference between using two types of methods as inputs through our approach. We have selected
 890 chromosome 3 for validation here because it is a gram as the input that there is no coplanar sample.
 891 Table 3 reports the quantitative stability metrics of the Gram-based embedding and the 3DMax
 892 reconstruction for the same chromosome. We observe that the Gram top eigenvalues are very close
 893 (spectral gap $\delta \approx 0$), making its orientation unstable. 3DMax, on the other hand, shows a much
 894 larger spectrum spread ($\delta \approx 18.6$) and correspondingly stable orientation.

895 **Table 3:** Stability metrics for Chromosome 3 (320 kb).

	Top-3 spectrum	Rotation stability
Gram	[0.00547, 0.00546, 0.00502], $\delta = 0.0000$	2.83
3DMax	[70.62, 52.04, 41.30], $\delta = 18.58$	4.49e-5

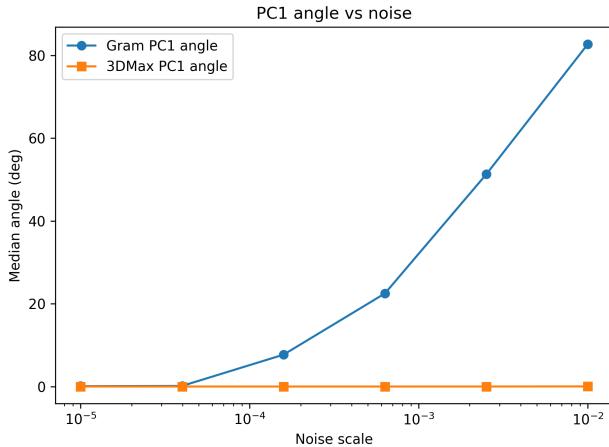
902 **Noise experiment.** Figure 5 shows the median angle of the first principal component (PC1) as a
 903 function of injected noise scale. The Gram embedding (blue) displays rapid growth of the PC1
 904 angle once noise is added, whereas the 3DMax embedding (orange) remains essentially unaffected.
 905 This is consistent with the theoretical bound above: a vanishing spectral gap makes the orientation
 906 of the Gram embedding highly sensitive to perturbations. These observations jointly validate our
 907 perturbation analysis: even when distances and neighborhood structure are perfectly preserved, a
 908 near-degenerate spectrum leads to unstable principal directions, whereas a large spectral gap confers
 909 rotational stability.

910 **IG-Gram experiment.** Table 4 illustrates the performance of Gram reconstruction as input into
 911 InertialGenome.

912 **Table 4:** dSCC and dRMSE of IG-Gram on different resolutions.

Method	320KB		160KB		80KB		40KB	
	dSCC↑	dRMSE↓	dSCC↑	dRMSE↓	dSCC↑	dRMSE↓	dSCC↑	dRMSE↓
IG-Gram	0.4770	0.3728	0.1079	0.4656	0.0524	0.4811	0.0083	0.5383

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933



934 Figure 5: Median PC1 angle versus noise scale for the Gram and 3DMax.
935

936 F RECONSTRUCTION PERFORMANCE ON B-LYMPHOCYTE CELL DATASET

937
938 Table 5 compares six methods on single-cell Hi-C data.
939

940 Table 5: Performance comparison of six methods on 3D chromosome structure reconstruction from
941 single-cell Hi-C data (B-Lymphocyte cell test set). Metrics report distance-based Spearman correlation
942 (dSCC \uparrow) and root mean square error (dRMSE \downarrow) at four resolutions. Best results in bold.
943

Method	1MB		500KB		250KB		100KB	
	dSCC \uparrow	dRMSE \downarrow						
3DMAX	0.9131	146.7464	0.8603	139.7606	0.8148	136.9766	0.6548	127.9119
LorDG	0.7462	107.7091	0.8103	117.5711	0.8316	112.0299	0.8395	98.9512
HiC-GNN	0.6778	0.2570	0.6457	0.2373	0.5827	0.1312	0.5334	0.2089
HiCEGNN	0.8847	0.0839	0.8068	0.0838	0.7530	0.0823	0.8017	0.0795
IG-3DMAX	0.9209	0.0822	0.9081	0.0777	0.8861	0.0593	0.8708	0.0790
IG-LorDG	0.8413	0.1114	0.8367	0.0979	0.8440	0.0675	0.7939	0.0867

954 G SUPPLEMENTARY FOR ABLATION STUDY

955 G.1 LOSS COMPONENTS

956 We analyze the contribution of the structural stability loss ($\mathcal{L}_{\text{struct}}$) versus the coordinate regression
957 loss ($\mathcal{L}_{\text{weighted_mse}}$) by varying their weighting ratio α/β in the total objective $\mathcal{L}_{\text{total}} = \alpha\mathcal{L}_{\text{struct}} +$
958 $\beta\mathcal{L}_{\text{weighted_mse}}$, where $\beta = 1 - \alpha$. As shown in Table 6, removing structural supervision entirely
959 ($\alpha = 0$) yields reasonable performance at coarse resolutions but leads to significant degradation
960 in dRMSE at fine scales (e.g., 40kb), indicating poor geometric consistency. In contrast, using
961 only structural loss ($\alpha = 1$) improves dRMSE at coarse resolutions but harms distance correlation
962 (dSCC). A certain proportion of structural regularization ($\alpha = 0.1, 0.5$) consistently achieves the
963 best trade-off across all resolutions, validating its role in enforcing biologically plausible 3D genome
964 topology without sacrificing coordinate fidelity.
965

966 G.2 COMPONENT ABLATION

967 Results using IG-3DMAX in Table 7 show: (1) Without inertial alignment, dRMSE increases across
968 all resolutions (e.g., from 0.1547 to 0.1641 at 320 kb), indicating its role in stabilizing global struc-
969 ture; (2) Removing RoPE consistently degrades both dSCC and dRMSE, confirming that relative
970 positional encoding is essential for structural fidelity; (3) Disabling the Nyström branch leads to the
971

972
 973 Table 6: Ablation study about IG-3DMAX on the loss weighting between structural stability and
 974 coordinate regression. Reported metrics: distance-based Spearman correlation (dSCC \uparrow) and root
 975 mean square error (dRMSE \downarrow) across four genomic resolutions. Best results per resolution are
 976 bolded.

977 978 979 980 981 982 983 984 985 986 987 988 989 990 991 992 993 994 995 996 997 998 999 1000 1001 1002 1003 1004 1005 1006 1007 1008 1009 1010 1011 1012 1013 1014 1015 1016 1017 1018 1019 1020 1021 1022 1023 1024 1025 Ratio (α/β)	320KB		160KB		80KB		40KB	
	dSCC \uparrow	dRMSE \downarrow						
0.0 / 1.0	0.9030	0.1728	0.8627	0.1935	0.7532	0.2152	0.7132	0.2410
0.1 / 0.9	0.9029	0.1696	0.8595	0.1848	0.7663	0.2197	0.7158	0.2407
0.5 / 0.5	0.9002	0.1671	0.8580	0.1879	0.7741	0.2266	0.7203	0.2445
1.0 / 0.0	0.8815	0.1453	0.8484	0.1775	0.7677	0.2297	0.7192	0.2788

largest performance drop at fine scales (e.g., dRMSE rises by 0.0114 at 40 kb), demonstrating its critical contribution to modeling long-range pairwise distances. The full model achieves the best trade-off between structural consistency and coordinate accuracy, outperforming all ablated variants at every resolution.

Table 7: Ablation study of key components in our model on single-cell Hi-C data (Frontal cortex test set). Reported metrics: distance-based Spearman correlation (dSCC \uparrow) and root mean square error (dRMSE \downarrow) across four genomic resolutions. Best results per resolution are bolded.

Model	320KB		160KB		80KB		40KB	
	dSCC \uparrow	dRMSE \downarrow						
Full (Ours)	0.9030	0.1547	0.8621	0.1809	0.7757	0.2035	0.7297	0.2382
w/o Inertial	0.9008	0.1641	0.8598	0.1869	0.7737	0.2185	0.7226	0.2385
w/o RoPE	0.8976	0.1613	0.8566	0.1894	0.7709	0.2229	0.7213	0.2454
w/o Nyström	0.9002	0.1659	0.8607	0.1998	0.7746	0.2218	0.7214	0.2496

H EXPERIMENTAL VALIDATION

H.1 VALIDATION VIA A/B COMPARTMENTALIZATION

To assess the biological plausibility of our predicted 3D chromosome structures, we validate them based on A/B compartment organization. In eukaryotic nuclei, chromosomes spatially segregate into two major compartments—A (active) and B (inactive)—with loci from the same compartment (A–A or B–B) tending to be spatially closer than loci from different compartments (A–B).

We follow the process of Lieberman-Aiden et al. (2009) to assign compartment labels. For each chromosome, we compute the Pearson correlation matrix of the normalized Hi-C contact matrix, then extract the first principal component (PC1). Loci with positive PC1 values are assigned to compartment A, and those with negative PC1 values to compartment B.

Distance computation and grouping. From the predicted 3D coordinates we compute pairwise Euclidean distances and categorize them into three groups:

- *intra-A*: distances between loci within compartment A;
- *intra-B*: distances between loci within compartment B;
- *inter-AB*: distances between loci from compartments A and B.

Statistical analysis and visualization. We compare the distributions of these three groups using box plots and permutation test. A valid 3D structure should exhibit significantly shorter distances for intra-compartment pairs (A–A and B–B) than for inter-compartment pairs (A–B), indicating that the predicted structures preserve the compartmental spatial organization.

H.2 VALIDATION VIA TAD DOMAIN CONSISTENCY

Topologically Associating Domains (TADs) are contiguous genomic regions within which loci tend to interact more frequently with each other than with loci outside the domain (Dixon et al., 2012).

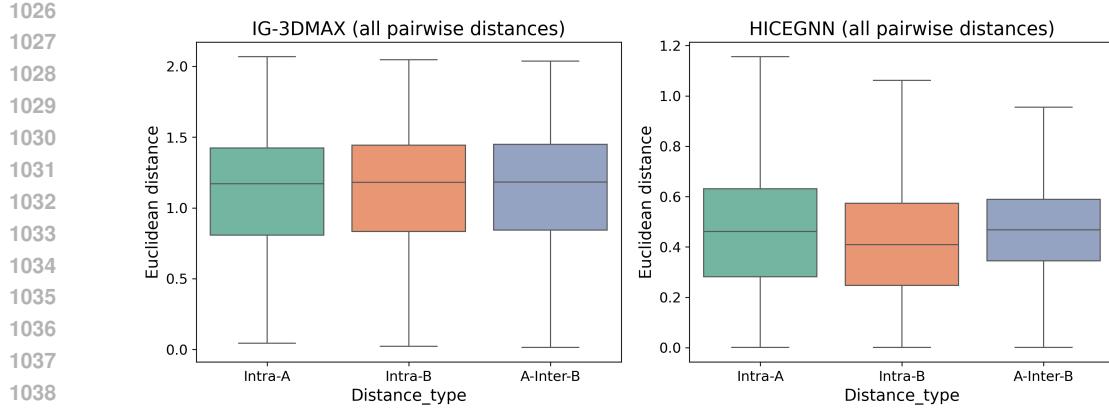


Figure 6: A/B compartment validation. Box plots show intra- and inter-compartment distances for IG-3DMDA (left) and HICEGNN (right).

Validating the biological plausibility of predicted 3D chromosome structures can therefore be performed by examining whether spatial distances among loci within the same TAD are significantly shorter than those between loci from different TADs.

TAD Boundary Acquisition TAD regions can be obtained from public annotations (e.g., ENCODE project or 3D Genome Browser) or identified directly from Hi-C contact matrices using established TAD callers such as: TopDom (Shin et al., 2016), Armatus (Filippova et al., 2014), and HiCExplorer (Wolff et al., 2020). In our workflow, TADs were identified using TopDom following the procedure in Serra et al. (2017).

Mapping 3D coordinates to TADs: Each bin in the 3D structure corresponds to a genomic interval based on Hi-C resolution (e.g., 250 kb). For a bin indexed by i , its genomic span is $[i \times 250 \text{ kb}, (i + 1) \times 250 \text{ kb}]$. We assign each 3D point to a TAD by checking whether its genomic start coordinate falls within the start and end of any TAD interval.

Evaluation Metric (Mean intra-/inter-TAD distance): For each TAD, we compute the average pairwise Euclidean distance D_{intra} among all loci inside the same TAD. We also compute the average distance D_{inter} between loci across different TADs. A well-structured 3D prediction should satisfy:

$$\text{mean}(D_{\text{intra}}) \ll \text{mean}(D_{\text{inter}}), \quad (37)$$

indicating spatial clustering of loci within the same TAD, consistent with known biological organization.

Statistical Testing We assessed whether intra-TAD distances are significantly smaller than inter-TAD distances using the Mann–Whitney U test. Table 8 summarizes the mean intra- and inter-TAD distances, intra/inter ratios, and corresponding p-values for IG-3DMDA and IG-LorDG reconstructions across chromosomes 4, 14, 18, and 20.

Table 8: TAD-based validation statistics for IG-3DMDA and HICEGNN reconstructions. Mean intra- and inter-TAD distances, intra/inter ratio, and Mann–Whitney U test p -values are reported for selected chromosomes.

Model	Chromosome	Intra-TAD	Inter-TAD	Ratio (Intra/Inter)	U test p -value
IG-3DMDA	Chr4	0.852	1.122	0.760	0
IG-3DMDA	Chr14	0.957	1.203	0.796	8.95e-93
IG-3DMDA	Chr18	0.882	1.083	0.814	1.55e-66
IG-3DMDA	Chr20	0.941	1.177	0.800	2.83e-47
HICEGNN	Chr4	0.408	0.442	0.925	2.61e-34
HICEGNN	Chr14	0.405	0.425	0.952	1.09e-4
HICEGNN	Chr18	0.380	0.416	0.914	3.16e-10
HICEGNN	Chr20	0.451	0.454	0.993	0.159

1080 These metrics together provide a quantitative assessment of whether the reconstructed 3D structures
 1081 preserve known domain-level chromatin organization. In particular, while IG-3DMAX consistently
 1082 shows significantly lower intra-TAD distances with low intra/inter ratios across all chromosomes,
 1083 HiCEGNN exhibits weaker separation and a non-significant result for chromosome 20, indicating
 1084 comparatively less biologically plausible domain organization in that case.
 1085

1086 H.3 FISH-BASED VALIDATION OF 3D STRUCTURES

1088 FISH experiments provide direct measurements of the spatial distances between genomic loci in the
 1089 nucleus. To validate our 3D chromatin structure predictions, we compared the predicted distances
 1090 between key regions—L1, L2, and L3—with the experimentally observed FISH distances reported
 1091 by Rao et al.(Rao et al., 2014). These regions were identified as loop anchors using HiCCUPS
 1092 analysis, with L1 and L2 forming a strong interaction peak, while L3 served as a non-interactive
 1093 control.

1094 We evaluated our model outputs at 250 kb resolution, where sufficient structural detail is preserved
 1095 to accurately localize these regions. For each chromosome, we computed the Euclidean distances
 1096 between L1–L2 and L2–L3 in the predicted 3D structures, along with their corresponding contact
 1097 probabilities from the original Hi-C data.

1098 As shown in Table 9, the predicted L1–L2 distances are consistently shorter than the L2–L3 distances
 1099 across all chromosomes, consistent with the experimental observations. Moreover, the contact prob-
 1100 abilities exhibit an inverse trend: higher values for L1–L2 compared to L2–L3, reflecting stronger
 1101 physical proximity. This agreement between predicted distances and Hi-C contact frequencies sup-
 1102 ports the biological plausibility of our modeled structures.

1103 Table 9: FISH validation results on GM12878 chromosomes 11, 14, and 17 at 250 kb resolution.
 1104 The table shows the predicted L1–L2 and L2–L3 distances, and the corresponding KR-normalized
 1105 contact probabilities from Hi-C data.

1107 Chromosome	1108 L1–L2 Distance	1108 L1–L2 Probability	1108 L2–L3 Distance	1108 L2–L3 Probability
1109 11	1109 0.8	1109 2.74×10^4	1109 3.3	1109 3.92×10^3
1110 14	1110 2.2	1110 9.10×10^3	1110 13.1	1110 2.07×10^3
1111 17	1111 2.3	1111 2.05×10^4	1111 3.7	1111 1.32×10^4

1112 The results demonstrate that our model captures the expected spatial organization: looped regions
 1113 (L1–L2) are closer in space and exhibit higher contact frequencies than non-looped regions (L2–L3).
 1114 This consistency with both FISH measurements and Hi-C data confirms that our method produces
 1115 biologically realistic chromatin structures.

1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133