



HAL
open science

Inference of 3D genome architecture by modeling overdispersion of Hi-C data

Nelle Varoquaux, William Noble, Jean-Philippe Vert

► **To cite this version:**

Nelle Varoquaux, William Noble, Jean-Philippe Vert. Inference of 3D genome architecture by modeling overdispersion of Hi-C data. *Bioinformatics*, 2023, 39 (1), 10.1093/bioinformatics/btac838 . hal-04234569

HAL Id: hal-04234569

<https://hal.science/hal-04234569>

Submitted on 10 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Genome analysis

Inference of 3D genome architecture by modeling overdispersion of Hi-C data

Nelle Varoquaux ^{1,*}, William S. Noble ^{2,3} and Jean-Philippe Vert^{4,5}

¹TIMC, Université Grenoble Alpes, CNRS, Grenoble INP, Grenoble 38000, France, ²Department of Genome Sciences, University of Washington, Seattle, WA 98195, USA, ³Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA 98195, USA, ⁴Brain Team, Google Research, Paris 75009, France and ⁵Centre for Computational Biology, MINES ParisTech, PSL University, Paris 75006, France

*To whom correspondence should be addressed.

Associate Editor: Valentina Boeva

Received on April 25, 2022; revised on November 16, 2022; editorial decision on November 29, 2022

Abstract

Motivation: We address the challenge of inferring a consensus 3D model of genome architecture from Hi-C data. Existing approaches most often rely on a two-step algorithm: first, convert the contact counts into distances, then optimize an objective function akin to multidimensional scaling (MDS) to infer a 3D model. Other approaches use a maximum likelihood approach, modeling the contact counts between two loci as a Poisson random variable whose intensity is a decreasing function of the distance between them. However, a Poisson model of contact counts implies that the variance of the data is equal to the mean, a relationship that is often too restrictive to properly model count data.

Results: We first confirm the presence of overdispersion in several real Hi-C datasets, and we show that the overdispersion arises even in simulated datasets. We then propose a new model, called Pastis-NB, where we replace the Poisson model of contact counts by a negative binomial one, which is parametrized by a mean and a separate dispersion parameter. The dispersion parameter allows the variance to be adjusted independently from the mean, thus better modeling overdispersed data. We compare the results of Pastis-NB to those of several previously published algorithms, both MDS-based and statistical methods. We show that the negative binomial inference yields more accurate structures on simulated data, and more robust structures than other models across real Hi-C replicates and across different resolutions.

Availability and implementation: A Python implementation of Pastis-NB is available at <https://github.com/hiclib/pastis> under the BSD license.

Contact: nelle.varoquaux@univ-grenoble-alpes.fr

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

DNA *in vivo* is folded in three dimensions, and this 3D structure plays an important role in many biological functions, including gene regulation, DNA replication, and DNA repair (De and Michor, 2011; Dixon *et al.*, 2012; Ryba *et al.*, 2010; Shen *et al.*, 2012). Chromosome conformation capture methods, coupled with next-generation sequencing, allow researchers to probe the 3D structure of chromosomes within the nucleus (Lieberman-Aiden *et al.*, 2009). These techniques, which we broadly refer to as ‘Hi-C,’ rely on cross-linking, digesting, ligating, and paired-end sequencing of DNA to identify physical interactions between pairs of loci. Hi-C techniques provide a genome-wide *contact map*, a matrix indicating the contact frequency between pairs of loci. This matrix can be used to analyze the 3D structure of the genome. However, despite extensive

research, inferring a 3D model from this contact map remains a fundamental problem.

Methods to infer the 3D structure of the genome broadly fall into two categories: *ensemble* approaches which infer a population of structures (Hu *et al.*, 2013; Kalhor *et al.*, 2012; Rousseau *et al.*, 2011; Zhu *et al.*, 2018), and *consensus* approaches which yield a single model that summarizes the contact count data (Ay *et al.*, 2014; Ben-Elazar *et al.*, 2013; Kapilevich *et al.*, 2019; Rieber and Mahony, 2017; Tanizawa *et al.*, 2010; Varoquaux *et al.*, 2014; Zhang *et al.*, 2013; 2019) (see Supplementary Table S1 for a more complete list of methods and their properties). The former approach is more biologically accurate because the population of models better reflects the diversity of structures present in a population of cells. However, interpretation of the resulting models is challenging, and one often has to fall back to a single structure or a few structures

that best represent the population of structures. Validation of ensemble models, even on simulated data, can also be challenging. On the other hand, a consensus model summarizes the hallmarks of genome architecture, can easily be visually inspected and analyzed, and can be integrated in a straightforward manner with other data sources, such as gene expression, methylation, and histone modifications, which are also ensemble based (Ay et al., 2014; Deng et al., 2015; Sexton et al., 2012; Varoquaux, 2019). In this work, therefore, we focus on inferring a consensus model of the 3D genome architecture.

Consensus approaches model chromosomes as chains of beads, minimizing a cost function that aims to produce a model as consistent with the data as possible. In addition, the optimization is sometimes constrained to include prior knowledge about the 3D structure: the size and shape of the nucleus, distance constraints between pairs of adjacent loci, etc. Some methods first convert contact counts into *wish distances*, either through a biophysically motivated counts-to-distances mapping or through *ad hoc* conversions. They then use multidimensional scaling (MDS) methods such as metric MDS (Duan et al., 2010; Rieber and Mahony, 2017; Tanizawa et al., 2010), weighted metric MDS (Ay et al., 2014; Zhang et al., 2019), non-metric MDS (Ben-Elazar et al., 2013) or classical MDS (most commonly known as PCA) (Kapilevich et al., 2019; Lesne et al., 2014; Li et al., 2018, 2020). These methods rely on arbitrary loss functions.

In previous work, we introduced a consensus method called ‘Pastis,’ based on a statistical model of contact count data, where the 3D structure is the latent variable and the inference of the consensus 3D model is formulated as a maximum likelihood problem (Varoquaux et al., 2014). A natural statistical model for count data is the Poisson distribution, which has a single parameter, the mean μ , from which all its other properties (including the variance) are derived. Pastis relies on such a Poisson model of the interaction frequencies, the intensity of which decreases with the increasing spatial distance between the pair of loci. We further extended this Poisson modeling to allow for the inference of diploid structures (Cauer et al., 2019). We refer to this method below as ‘Pastis-PM’ (‘PM’ standing for ‘Poisson model’).

However, a PM of contact counts implies that the variance of the data is equal to the mean, an assumption that is sometimes too restrictive to properly model the data. For instance, Nagalakshmi et al. (2008) and Robinson and Smyth (2007) show that for RNA-seq, this assumption is not justified, because the variance in the data is larger than the mean, leading to an overdispersion problem. Overdispersion typically occurs when observations deviate from the hypothesized structure of the data (e.g. dependency between observations, contagion, a cluster structure, and heterogeneity) (Xekalaki, 2014). To alleviate the overdispersion, Robinson et al. (2010) suggest modeling RNA-seq data with a negative binomial distribution, which is parameterized with two parameters: the mean μ and the variance σ . Modeling Hi-C contact count data with a negative binomial model is not new. Jin et al. (2013) and Carty et al. (2017) used this approach to assign statistical confidence estimates to observed contacts, Hu et al. (2012) to normalize the data, and Lévy-Leduc et al. (2014) to find topologically associated domains.

In this work, we explore methods that apply a similar generalization—from Poisson to negative binomial—in the context of a model for inferring genome 3D structure from a Hi-C contact map. We first confirm the overdispersion on a wide variety of Hi-C datasets, from very small (*Saccharomyces cerevisiae*) to large genomes (human). We then compare our method based on a negative binomial model for Hi-C count data, which we call Pastis-NB, to MDS-based methods (chromSDE, ShRec3D, Pastis-MDS, miniMDS, and SuperRec) and to the Poisson model-based Pastis-PM. We first demonstrate that Pastis-NB recovers the most accurate results, in particular in low coverage settings. We then study how well the different methods perform when provided with an incorrect mapping between contact counts and Euclidean distances, a setting where Pastis-NB also outperforms other methods. Finally, we show that the Pastis-NB model yields more stable structures across Hi-C replicates and across resolutions.

2 Approach

We model each chromosome as a series of n beads, where each bead (or *locus*) corresponds to a specific genomic window. We aim to infer the coordinate matrix $\mathbf{X} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^{3 \times n}$, where x_i corresponds to the 3D coordinates of bead i . We denote by $d_{ij} = \|x_i - x_j\|$ the Euclidean distance between beads i and j . Hi-C contact counts can be summarized as an n -by- n matrix c in which rows and columns correspond to loci, and each entry c_{ij} is an integer, called a *contact count*, corresponding to the number of times loci i and j are observed in contact. This matrix is by construction square and symmetric. Since raw contact counts are known to be biased by locus-specific multiplicative factors, we apply ICE normalization (Imakaev et al., 2012) to estimate a bias vector $b = (b_1, \dots, b_n) \in \mathbb{R}^n$, where b_i is the bias factor for locus i . The normalized contact count matrix c^N is then defined as the matrix of normalized counts $c_{ij}^N = c_{ij}/(b_i b_j)$. See Supplementary Material S4.1 for more details.

2.1 Statistical model

We model the raw contact counts c_{ij} as realizations of independent negative binomial random variables

$$C_{ij} \sim \text{NB}(r_{ij}, \mu_{ij}),$$

where r_{ij} is the dispersion parameter and μ_{ij} is the mean of the negative binomial distribution between loci i and j . Like in the Pastis model, we parameterize the mean count value μ_{ij} as a decreasing function of the distance between beads i and j : $\mu_{ij} = \beta b_i b_j d_{ij}^\alpha$, with parameters $\beta > 0$ and $\alpha < 0$. β can be thought of as a scaling factor—the higher the coverage of the dataset, the higher it is—while α characterizes how the frequency of contacts decreases with the distance. Note that this relationship is ill-posed, in the sense that the relationship between the mean of the distribution μ , the scaling factor β , and the Euclidean distances d is not unique. In particular, one can choose to set β to 1 in order to infer the 3D coordinates, and then rescale the inferred structure to reflect prior knowledge about the size of the nucleus. We thus drop the scaling parameter β in the rest of the derivations. In addition, we parameterize the dispersion as $r_{ij} = b_i b_j r$, where $r \geq 0$ accounts for overdispersion.

The probability mass function can thus be written as:

$$\Pr(C_{ij} = c_{ij}) = \frac{\Gamma(c_{ij} + r_{ij})}{\Gamma(c_{ij} + 1)\Gamma(r_{ij})} \left(\frac{\mu_{ij}}{r_{ij} + \mu_{ij}}\right)^{c_{ij}} \left(\frac{r_{ij}}{r_{ij} + \mu_{ij}}\right)^{r_{ij}}.$$

It is well known that the variance of C_{ij} satisfies

$$\sigma_{ij}^2 = \mu_{ij} + \frac{\mu_{ij}^2}{r_{ij}} = b_i b_j \left(d_{ij}^\alpha + \frac{d_{ij}^{2\alpha}}{r}\right),$$

and that when $r_{ij} \rightarrow \infty$ the negative binomial distribution tends to a Poisson distribution with intensity parameter μ_{ij} . The negative binomial is thus a generalization of the Poisson distribution where the variance of the data can exceed the mean, as controlled by the dispersion parameter (Nikoloulopoulos and Karlis, 2008).

2.2 Estimating the dispersion parameters r_{ij}

To estimate the dispersion parameters r_{ij} , we leverage the property that the variance of an $\text{NB}(\rho, \mu)$ random variable is $\sigma^2 = \mu + \mu^2/\rho$; therefore, if we know μ and σ^2 , we deduce the dispersion as $\rho = \mu^2/(\sigma^2 - \mu)$. This implies that if a relationship $\sigma^2(\mu)$ between the mean and variance of the NB distribution is known or assumed, then the dispersion parameter ρ is also known as a function of μ as $\rho(\mu) = \mu^2/(\sigma^2(\mu) - \mu)$. We thus focus on estimating the variance as a function of the mean of the contact counts in order to infer the dispersion parameter r .

In the case of RNA-seq, the function $\sigma^2(\mu)$ is usually estimated by fitting a weighted least squares (Behr et al., 2013), fitting a loess (Anders and Huber, 2010), or by maximizing a likelihood on empirical means and variances estimated for each gene using replicates (Robinson and Smyth, 2008). In the case of Hi-C, relying on

biological replicates is not a viable option because most studies perform at most two biological replicates, rendering the estimation of the mean and variance impossible. Instead, we estimate $\sigma^2(\mu)$ by introducing additional modeling assumptions that capture properties intrinsic to Hi-C data and genome architecture. We propose a two-step method to estimate this function. First, we compute for each genomic distance l the empirical mean and variance of the normalized contact counts:

$$\begin{aligned}\tilde{\mu}_l &= \frac{1}{N_l} \sum_{(i,j):|i-j|=l} \frac{c_{ij}}{b_i b_j}, \\ \tilde{\omega}_l &= \frac{1}{N_l - 1} \sum_{(i,j):|i-j|=l} \left(\frac{c_{ij}}{b_i b_j} - \tilde{\mu}_l \right)^2,\end{aligned}$$

where N_l is the number of (i, j) pairs with $|i - j| = l$. As shown by Yu *et al.* (2013), $\tilde{\omega}_l$ is a biased estimator of the variance and can be corrected to be unbiased as follows:

$$\tilde{\sigma}_l^2 = \frac{\tilde{\omega}_l}{\frac{1}{N_l} \sum_{(i,j):|i-j|=l} \frac{1}{b_i b_j}}.$$

Note that in our model, estimating the mean and variance from empirical normalized counts at a given genomic distance amounts to assuming that the mean normalized count, i.e. d_{ij}^z , is constant for beads (i, j) at a given genomic distance from each other. This is obviously an assumption that we only use to get an estimate of the overdispersion parameter, but that we relax later when we optimize the 3D coordinates of each bead without constraint on their pairwise distances.

Because the empirical mean and variance may not be reliable for very long genomic distances, where the number of contact counts is small, we only compute $\tilde{\mu}_l$ and $\tilde{\sigma}_l^2$ for $l < \frac{2}{3} l_{\max}$, where l_{\max} is the maximum distance between loci. We also discard genomic distances with an empirical mean or variance equal to 0.

We then proceed to estimate the dispersion parameter r in two steps. First, from the set of $(\tilde{\mu}_l, \tilde{\sigma}_l^2)$ pairs, we then estimate the dispersion \tilde{r}_l for all genomic distance as follows:

$$\tilde{r}_l = \frac{\tilde{\mu}_l^2}{(\tilde{\sigma}_l^2 - \tilde{\mu}_l)}.$$

Second, from the estimates \tilde{r}_l for each genomic distance, we proceed to estimating the final dispersion parameter r by taking the weighted average:

$$r = \max\left(0, \frac{\sum_l w_l \tilde{r}_l}{\sum_l w_l}\right),$$

where w_l the number of data points used in the estimation of $\tilde{\mu}_l$ and $\tilde{\sigma}_l^2$. Thus, more weight is given to short genomic distances than long genomic distances.

2.3 Estimating the 3D structure

In summary, our proposed model has three main components: (i) we model contact counts using negative binomial distributions parameterized by the mean and the dispersion parameter $C_{ij} \sim NB(r_{ij}, \mu_{ij})$; (ii) we parameterize the mean as a function of the structure $\mu_{ij} = b_i b_j d_{ij}^z$; and (iii) we model the dispersion parameter as $r_{ij} = b_i b_j r$ and provide a method to estimate r . By combining these three components, we can write the probability of each observation:

$$\begin{aligned}P(C_{ij} = c_{ij}) &= \frac{\Gamma(c_{ij} + b_i b_j r)}{\Gamma(c_{ij} + 1) \Gamma(b_i b_j r)} \\ &\quad \times \left(\frac{d_{ij}^z}{r + d_{ij}^z} \right)^{c_{ij}} \left(\frac{r}{r + d_{ij}^z} \right)^{b_i b_j r},\end{aligned}$$

which depends on the 3D structure \mathbf{X} through the pairwise distances d_{ij} , and which also depends on the count-to-distance mapping parameter α . We then propose to jointly infer both the 3D structure \mathbf{X} and the α parameter by maximizing the likelihood:

$$\max_{\mathbf{X}} \mathcal{L}(\mathbf{X}, \alpha) = \sum_{(i,j)} \log(P(C_{ij} = c_{ij})). \quad (1)$$

Given an observed Hi-C contact map, we solve the optimization problem of Equation (1) using the L-BFGS algorithm (Byrd *et al.*, 1995; Liu and Nocedal, 1989) from the `scipy` toolbox. The optimization being non-convex, we solve it five times with random initialization to find local optima and return the solution with the highest log-likelihood. Note that the problem in \mathbf{X} is ill-posed, in the sense that the solution is defined relative to a rotation factor and a translation factor. Note also that Equation (1) is a sum over all pairs of loci (i, j) , and may include zero counts $c_{ij} = 0$ which contribute to the likelihood. We illustrate below the benefits of keeping zero counts in the objective functions and not filtering them.

3 Results

We perform a series of experiments to assess the accuracy of the different methods on simulated data and the robustness of the methods on real data. Specifically, we perform experiments on simulated data to assess the accuracy of the inferred 3D models when varying (i) the coverage of the datasets; (ii) the overdispersion of the counts; and (iii) the counts-to-distance mapping parameter. We then assess the stability of the methods on real data (i) on two different biological replicates; (ii) across different resolution; (iii) when subsampling the data; (iv) on a multi-chromosome dataset, by varying the number of chromosomes included in inference. We present a summary of all results in Table 1. See Supplementary Methods for more detailed on data simulation, pre-processing and validation metrics used.

3.1 Real and simulated Hi-C data are highly overdispersed

Before diving into the comparison of the different models and algorithms, we first investigate the extent to which existing Hi-C datasets show evidence of overdispersion. Because it is rare to have several samples of the same cell line, we use the method described in Section 2 to check for the presence or absence of overdispersion in the normalized Hi-C data by estimating the mean and variance for each genomic distance. We plot in Figure 1, for different species (chr10 of the KBM7 human cell line at 100 kb, *Drosophila melanogaster* at 10 kb, *Arabidopsis thaliana* at 40 kb and *S.cerevisiae* at 10 kb), the mean versus variance relationship of normalized contact counts. Each dot in the plot corresponds to a particular genomic distance. In each case, we observe strong overdispersion (i.e. points above the diagonal), supporting the idea of modeling Hi-C count data with overdispersed models. See Supplementary Figure S1 for an estimation of the dispersion parameter per genomic distance.

We then repeat this experiment on simulated data by generating 50 000 structures using a previously described volume exclusion model of the budding yeast (Tjong *et al.*, 2012). From this population of structures, we create a contact count matrix at 3 kb resolution, assuming that loci closer than 40 nm come into contact. This simulated contact count map has been shown to be highly correlated with experimental Hi-C data (Tjong *et al.*, 2012). The resulting dataset (purple series in Fig. 1) displays the same overdispersion as the data. We thus conclude that the overdispersion is an inherent property of Hi-C data and not an experimental artifact. We hypothesize that the overdispersion arises due to the large variety of different structures present in a single Hi-C experiment.

3.2 Patis-NB is accurate and robust on simulated data

Next, we use simulated data to compare our approach, which we refer to as Patis-NB, with seven different algorithms. Patis-MDS is a weighted metric MDS method that attempts to place the beads such that the distance between each pair matches as closely as possible the wish distances derived from contact counts (Varoquaux *et al.*, 2014). The miniMDS method improves upon weighted metric MDS by progressively solving high-resolution overlapping ‘local’ structures and assembling them using low-resolution whole chromosome inferred structures (Rieber and Mahony, 2017). This strategy yields

Table 1. Summary of the results

	Simulated data			Real data				
	Coverage (small β)	Dispersion (small γ)	Counts-to-maps (small α)	Replicates (LR)	Replicates (HR)	Resolution	Downsampling	Multi-chrom.
Pastis-MDS	.	✓✓
miniMDS	✓✓	.	✓✓	✓	✓✓	✓	✓✓	✓
ShRec3D	✓✓	.	✓	✓	✓✓	✓	✓✓	✓
SuperRec
ChromSDE	.	.	.	✓✓✓	✓	.	.	NA
ShNeigh	.	.	.	✓✓✓	✓	.	.	NA
Pastis-PM	✓	✓	✓✓✓
Pastis-NB	✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓

Note: We summarize here all the experiments performed and a synopsis of the main results. The table displays the top three methods for each experiment (✓✓✓ for the best method, ✓✓ for the second best, ✓ for the third best). Pastis-NB performs better than all methods in most experiments.

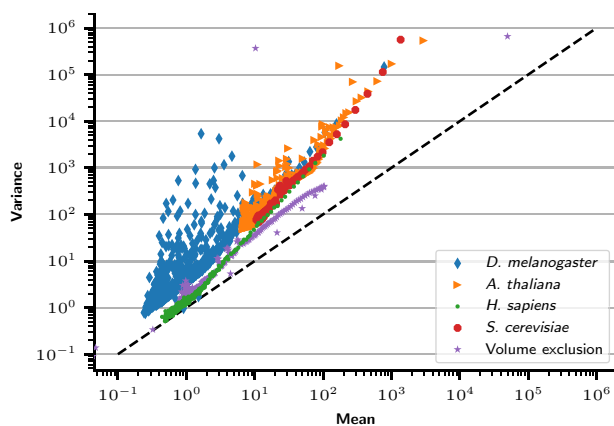


Fig. 1. Mean and variance of contact counts in different Hi-C datasets. Each point represents a given genomic distance in one dataset, where we estimate the mean and variance of contact counts. The dashed line corresponds to the relationship assumed by the PM: $\sigma^2 = \mu$. The ‘Volume exclusion’ dataset is simulated following a previously described model (Tjong et al., 2012)

improvement both in terms of computation time and stability. ChromSDE is a variant of metric MDS that penalizes non-interacting pairs of beads to keep them far away from one another and optimizes the counts-to-distances mapping coefficient with a golden search (Zhang et al., 2013). ShRec3D is a two-step method that first derives distances from contact counts using a shortest path algorithm and then applies classical MDS (Lesne et al., 2014). SuperRec improves upon ShRec3D in two ways: (i) optimizing the counts-to-distances mapping coefficient with a golden search, similar to chromSDE; (ii) using an iterative optimization strategy by inferring overlapping local structures and then combining them to infer a full whole-genome chromosome. ShNeigh combines the classical MDS approach with a local Gaussian model of neighboring beads (Li et al., 2020). Finally, Pastis-PM models contact counts as Poisson random variables with the 3D structure as a latent variable and casts the inference problem as a likelihood maximization, optimizing jointly the structure and the parameters of the counts-to-distance function. The eight methods fall into two categories: chromSDE, SuperRec, Pastis-PM, and Pastis-NB are *non-metric* because they fit a parametric curve to estimate a count-to-distance mapping from the data, whereas Pastis-MDS, miniMDS, ShRec3D, and ShNeigh are *metric* because they do not do this fitting. In our experiments, each software package is used with default parameters.

3.2.1 Robustness to coverage and dispersion

We first assess how well the different methods reconstruct a known 3D structure from simulated data at different coverage and

dispersion levels. High coverage typically corresponds to a high signal-to-noise ratio, whereas low coverage yields sparse, low signal-to-noise ratio matrices. Similarly, when the dispersion parameter tends to infinity, the negative binomial distribution (by definition overdispersed) tends to a Poisson with lower variance. Thus, the lower the dispersion parameter, the noisier the dataset. All methods should therefore perform better as the dispersion parameter (γ in our setting) and the coverage increase.

In this first series of experiments, we provide the correct count-to-distance or distance-to-count transfer functions to the metric methods, who need it. In this setting, for infinite coverage and infinite dispersion parameter, all methods should consistently estimate the correct structure, at least if they manage to converge to the global optimum of their objective function.

We first plot the average root mean squared deviation (RMSD) error between the true and predicted structures, as a function of coverage (Fig. 2A) for the datasets with varying coverage. Strikingly, ShRec3D’s, SuperRec’s, ShNeigh’s and Pastis-NB’s results are extremely stable to coverage, while the four other methods see their performance decrease with decreasing coverage (only slightly for miniMDS). ShRec3D, ShNeigh, and SuperRec perform relatively well when the coverage is low but poorly in the high coverage setting. All the other methods perform similarly well for high coverage, with Pastis-PM and Pastis-NB achieving the lowest RMSD at high coverage, but exhibit strong differences in the low coverage setting. In the low coverage setting, Pastis-NB remains extremely good, with barely any decrease in performance even at the lowest coverage, while Pastis-MDS’ performance degrades quickly with less than 70% of the coverage, and similarly Pastis-PM and chromSDE see their performance deteriorate with less than 40% of the coverage. The miniMDS method outperforms greatly Pastis-MDS in the low coverage setting despite the two methods being very related, confirming that classical metric MDS methods are dominated by long range (and thus noisy) interactions. With the best RMSD error among all methods for all coverage, Pastis-NB is the clear winner in this experiment.

We then plot the average RMSD as a function of the dispersion tuning parameter γ (Fig. 2B) for the second set of simulated contact maps with varying dispersion. As expected, we observe that all methods tend to perform better when γ increases (corresponding to less overdispersion) and have poorer performance when the data are too overdispersed. ShRec3D’s and ShNeigh’s results tend to be stable to changes in dispersion, but worse than other methods for large γ ’s. All methods perform poorly in the highest dispersion setting. ChromSDE performs relatively poorly in the medium to high dispersion setting. SuperRec, ShRec3D, and ShNeigh perform relatively poorly in all setting. Pastis-NB has again the best performance across all dispersion values, although the difference relative to other methods, particularly miniMDS, Pastis-MDS, and Pastis-PM, is small.

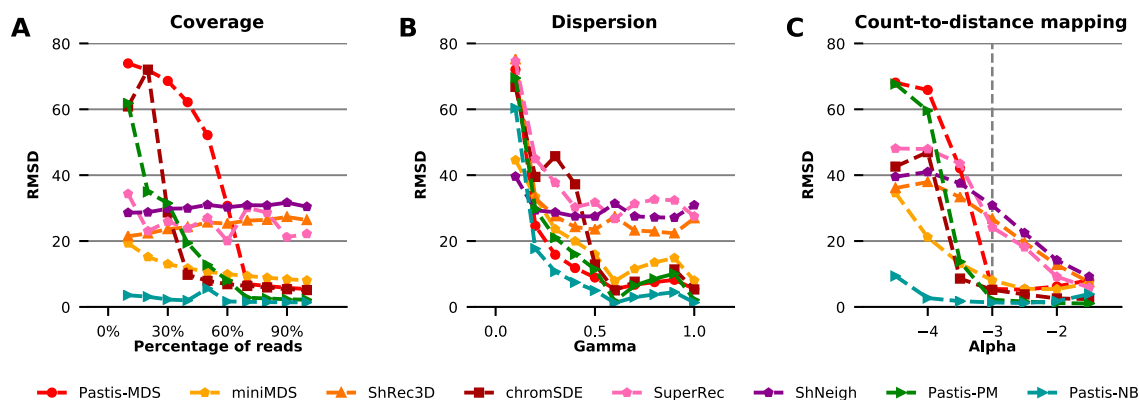


Fig. 2. Performance evaluation on simulated data. Each plot shows the mean RMSD error (over 10 random simulations with different random seed) between the predicted structure and the true structure, for eight different methods, when one parameter of the simulation is varied. (A) The parameter β is varied such that the coverage is 10–100% of the original dataset. (B) The parameter γ , which controls overdispersion, is varied between 0.1 and 1. Smaller values correspond to more overdispersion. (C) The parameter α , which controls the count-to-distance mapping coefficient, is varied between -4.5 and -1.5 .

3.2.2 Robustness to incorrect parameter estimation

First, we explore the robustness of Pastis-NB with respect to the dispersion parameter estimation. For that purpose, we re-run the inference on the simulated datasets with a varying coverage estimation using three different estimated dispersion parameters: our original estimated dispersion parameter \hat{r} , $\hat{r}/10$ (i.e. 10-fold underestimation of the dispersion of the datasets), and $\hat{r} \times 10$ (i.e. 10-fold overestimation of the dispersion). We then compare the results to the second-best performing method on this dataset, miniMDS. [Supplementary Figure S2](#) shows that the performance of Pastis-NB is barely affected by under- or overestimation of the dispersion parameter, and that it outperforms miniMDS on a wide range of dispersion parameter estimates. The fact that the results are slightly better with a larger dispersion parameter suggests that we tend to underestimate the dispersion parameter of the negative binomial model for the counts (or, equivalently, overestimate the variance of the counts as a function of their means). One reason for this systemic bias could be that we bin together all counts for pairs of loci at a given genomic distance in order to estimate the mean and variance at each genomic distance, from which we estimate the dispersion parameter of our model. Since those pairs of loci at a given genomic distance are not exactly at the same spatial distance from each other, the count variance we estimate is due both to the overdispersion of the counts for a given spatial distance (which we want to estimate for our model), and to the variation in the spatial distances of loci at a given genomic distance (which is a nuisance parameter here). It is therefore likely that we overestimate the variance of the counts due to the overdispersion only, and that we therefore underestimate the dispersion parameter. A possibility to improve dispersion parameter estimation would be to estimate it from biological replicates instead of a single replicate.

In addition to demonstrating robustness to the dispersion parameter estimation, these results also show that the performance of Pastis-NB is barely affected by the *ad hoc* choice we made to estimate the constant dispersion \hat{r} from the genomic distance-dependent estimates \hat{r}_l only for $l \leq 2/3l_{\max}$, since taking a larger threshold up to $l \leq 0.99l_{\max}$ impacts the \hat{r} estimate by less than 10% on all datasets ([Supplementary Table S3](#)).

Second, we compare the algorithms on datasets with varying counts-to-distances mappings. Metric methods (Pastis-MDS, miniMDS, and ShRec3D) require as input a count-to-distance transfer function. While Pastis-MDS and miniMDS rely on ideal physical laws to define this mapping, ShRec3D uses *ad hoc* conversion of counts into physical distances. However, DNA may not follow the ideal properties of polymers underlying the default transfer function; thus, structures inferred from these methods may diverge from the correct ones. Our goal here is to assess how well the different methods perform when the transfer function is misspecified. We expect non-metric methods to perform better on these datasets, because

they should be able to adapt the transfer function to best fit the data.

[Figure 2C](#) shows the average RMSD error of each algorithm as a function of the α parameter used to simulate the data. It is worth noting that the lower the α parameter is, the noisier the simulated contact map is: a low α parameter indeed results in a contact count map with very few long range interaction counts.

Pastis-NB works well across all values of α , exhibiting a striking difference from the rest of the methods for $\alpha \leq -3$. Notably, all methods perform much better for high α than for low α . This phenomenon can be explained by the properties of the contact count maps in this setting: low α values in the count-to-distance function lead to abrupt changes in the probability of seeing contact counts between small and large distances, whereas high α values yield a much more uniform expected contact counts map. Thus, for identical coverage, low α datasets are much sparser than high α datasets. In short, despite identical coverage in all datasets, the signal-to-noise ratio varies strongly with α , thereby leading to much better overall performance for low α both for metric and non-metric methods, even when the transfer function is misspecified. Pastis-PM works very poorly in the low α setting, probably because the data is much more overdispersed in this setting than in the high α setting. Once again, miniMDS works much better than Pastis-MDS, underlining the benefits of only considering interactions from neighboring loci for precise local structure inference.

3.3 Pastis-NB yields stable and consistent structures on real Hi-C data

We now test the different methods on real Hi-C data. Since in this case the true consensus structure is unknown, we investigate instead the behaviors of the different methods in terms of their ability to infer consistent structures from replicate datasets and across resolutions.

3.3.1 Pastis-NB shows increased stability when performing the inference without filtering zero counts

Before delving into a detailed comparison of Pastis-NB with other methods on various tasks, we first illustrate in this section the benefits achieved by including zero-valued counts in the model inferred by Pastis-NB. This is to be contrasted with many previously published methods that perform 3D structure inference after using only a subset of the data, and in particular disregarding zero counts. For example, [Tanizawa et al. \(2010\)](#) and [citeduan:3D \(Duan et al., 2010\)](#) consider only the top 2% significant contact counts, whereas [Varoquaux et al. \(2014\)](#) exclude zero contact counts from the inference. Furthermore, many MDS-based methods require a transformation of contact counts into distances: this is often based on a power-law relationship with a negative coefficient and is thus

undefined for zero contact counts. As a result, such methods include zero contact counts either through an *ad hoc* penalization term on non-interacting beads or by converting zero counts to an *ad hoc* distance (e.g. the largest distance obtained on non-zero counts or using prior knowledge of the structure). In contrast, Pastis-NB formulates the inference in a way that naturally includes zero contact counts.

To assess the impact of zero counts information, we compare the default Pastis-NB model, which takes into account zero counts in the likelihood objective function (1), with a variant where we only retain non-zero counts. We assess the stability of the inferred structures across biological replicates when using all counts versus only non-zero counts. To do so, we run the inference on the whole sets of contact counts versus the filtered one, on both replicates 75 and 76 of KBM7, for contact count matrices at five different resolutions. We thus infer two structures X^1 and X^2 for each autosomal chromosome, one for each replicate. We then rescale them such that the 99% of the beads of the structures fit in a sphere of a predefined diameter, and we compute two RMSD measures: $\text{RMSD}(X^1, X^2)$ and $\text{RMSD}(X^2, X^1)$ and report the average value. We also compute the Spearman correlations between the distance matrices of the two structures. Those two measures assess the stability of the inference. Recall that the different methods cannot optimize the coordinates of beads that have zero contact counts. Thus, before computing the correlation, we filter out from both structures any beads that have zero contact counts in either dataset.

We then compare the stability of the inference between the two approaches (filtered versus unfiltered zero counts). Figure 3A (resp. Supplementary Fig. S3) shows, for each approach, the distribution of RMSD (resp. Spearman correlation) values across the 22 chromosomes and five resolutions. We clearly see that keeping all zero counts leads to significantly larger correlations, hence more stable structures across biological replicates.

3.3.2 Stability to replicates

Replicates, which involve multiple runs of the same experiment performed on similar samples with the same experimental settings, are typically carried out to assess variability of the results. Because the underlying 3D model should not change in this case, we compare the results of the inference of the different algorithms on two replicates of the nearly haploid human KBM7 cell line. Similar to what we did in the previous section, we infer two structures on the two replicates of interest for each autosomal chromosome. We compute the RMSD between each pairs of structures as described above, as well as the Spearman correlation of the distances.

Not surprisingly, the stability of the inference across replicates decreases as the resolution grows in all methods. Table 2 shows the average correlation and RMSD reached by each method at each resolution. At low resolution (1 Mb and 500 kb), chromSDE

performs the best both in terms of correlation and in terms of RMSD between replicates (except for RMSD at 1 Mb, where it is outperformed by SuperRec), although all methods perform well in that setting, with correlations ranging between 0.95 and 0.99. At high resolution (250, 100, and 50 kb), Pastis-NB performs the best, both in correlation and in RMSD, despite the non-convex nature of the optimization problem solved. It is remarkable that even at 50 kb, Pastis-NB reaches a correlation of 0.96 between replicates, while the second best method (ShNeigh) see its correlation decrease to 0.92 and RMSD to 13.1, and chromSDE, which is the most stable at low resolution, only obtains a correlation of 0.62 and an RMSD of 29.66.

3.3.3 Stability to resolution

Zhang et al. (2012) show that the mapping from contact counts to physical distance differs from one resolution to another, underscoring the importance of good parameter estimation. To study the stability of the structure inference methods to changes in resolution, we compute the RMSD and Spearman correlation between pairs of structures inferred at different resolutions (1 Mb, 500 kb, 250 kb, 100 kb, and 50 kb). Each inferred structure is rescaled such that all beads fit in a nucleus of size 100. To compare two structures at different resolutions, we downsample the structure of the highest resolution by averaging its coordinates until it is of the same resolution as the other one. We then compute the RMSD and Spearman correlation between the two structures of the same size.

Results of this experiment (Table 3) show that Pastis-NB is more stable to resolution changes than other methods, both in terms of RMSD and in terms of correlation. The second best method is ShNeigh, with stability measures almost as good as Pastis-NB. Surprisingly, the inferred models are visually extremely different (see Supplementary Fig. S6). Interestingly, the two subsequent contenders in terms of correlation are miniMDS and SuperRec, two methods that implement an iterative strategy of optimization, thus favoring close-range interactions. The poor performance of SuperRec in terms of RMSD is due to the presence of outlier predictions on a small number of reads.

3.3.4 Stability to coverage

We then study the stability of the structure inference methods to coverage. To do so, we downsample the 100 kb contact count matrices between 10% and 90% of the original coverage. We perform inference from these downsampled contact maps and compute the Spearman correlation between Euclidean distances of the obtained structures and the structures inferred on the full matrix. Results of this experiment (Fig. 3B, Supplementary Fig. S4) show that all methods tend to see the correlation decrease with the downsampling, as expected. While ShRec3D and chromSDE yield high correlations at

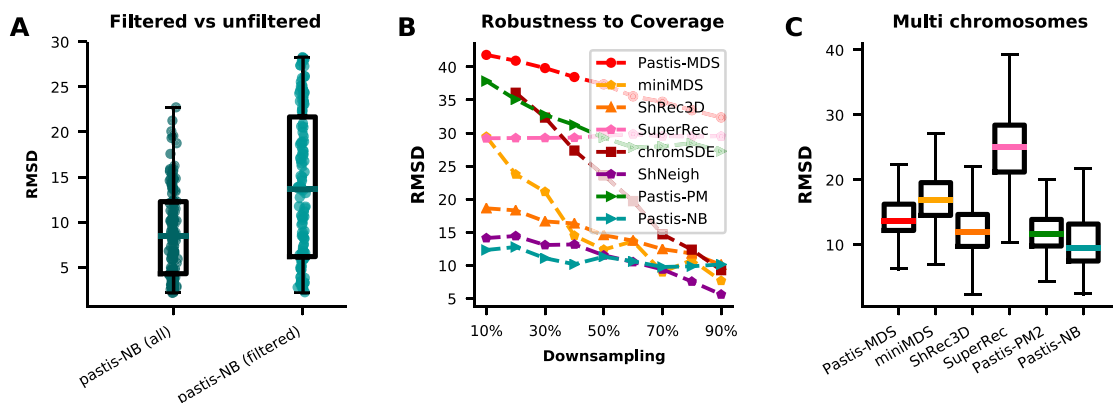


Fig. 3. Stability results on human single chromosomes Rao et al. (2014) and yeast whole genome Duan et al. (2010). (A) The RMSD between pairs of structures inferred by Pastis-NB using all contact counts ('all') or excluding zero contact counts ('filtered'). (B) Inference is performed on a downsampled contact map and the resulting structure is compared to the structure obtained using the whole dataset. ChromSDE fails to infer a structure with the correct number of bins on the datasets downsampled to 10% of the original coverage: results for chromSDE at 10% are thus not displayed. (C) The RMSD between pairs of chromosomes inferred by subsampling chromosomes from a *S.cerevisiae* whole genome dataset Duan et al. (2010)

Table 2. Stability across replicates

	pastis-MDS		miniMDS		ShRec3D		SuperRec		chromSDE		ShNeigh		pastis-PM2		pastis-NB	
	RMSD	Corr.	RMSD	Corr.	RMSD	Corr.	RMSD	Corr.	RMSD	Corr.	RMSD	Corr.	RMSD	Corr.	RMSD	Corr.
1mb	10.9	0.95	14.4	0.93	6.9	0.97	4.8	0.97	5.5	0.99	6.1	0.98	11.4	0.96	10.1	0.96
500kb	9.6	0.97	12.4	0.94	10.3	0.96	20.0	0.97	7.8	0.98	9.6	0.96	12.3	0.95	8.1	0.98
250kb	13.2	0.92	11.3	0.96	13.7	0.93	19.2	0.97	8.9	0.97	12.6	0.94	17.8	0.89	10.5	0.97
100kb	28.6	0.59	15.0	0.93	16.3	0.89	50.3	0.88	16.2	0.87	13.3	0.92	24.7	0.73	11.9	0.96
50kb	38.6	0.33	21.1	0.84	17.0	0.87	66.7	0.81	29.7	0.62	13.2	0.92	31.3	0.56	11.7	0.96

Note: The table shows the average RMSD and Spearman correlation between structures inferred on biological replicates on 22 autosomes of the KBM7 cell line at varying resolutions. In bold is the best correlation per row.

Table 3. Stability across resolution

	Pastis-MDS	miniMDS	ShRec3D	SuperRec	chromSDE	ShNeigh	Pastis-PM2	Pastis-NB
RMSD	29.5	20.5	19.6	38.1	23.0	17.6	24.6	17.0
Corr.	0.48	0.81	0.79	0.81	0.65	0.83	0.66	0.85

Note: The table lists the average Spearman correlation between pairs of structures for replicate 75 at different resolutions. In bold are the lowest average RMSD and highest average Spearman correlation. These values were computed on the KBM7 nearly haploid human cell line (Rao *et al.*, 2014).

high coverage, the correlation decreases sharply for chromSDE, reaching ~ 0.3 at 10% downsampling, and a bit less sharply for ShRec3D, which reaches ~ 0.8 at 10% downsampling. Pastis-PM and Pastis-MDS have the worst correlation at high coverage, and see their correlation decrease sharply with coverage. Pastis-NB and ShNeigh stand out as the methods with the largest correlation at high coverage but also as the methods that witness barely any decrease in correlation when coverage decreases.

3.3.5 Whole genome inference

Finally, we consider the harder task of whole genome inference, rather than inferring structures separately per chromosome. When tackling whole genome inference, a new problem arises for non-metric methods: the *inter*-chromosomal contact counts dominate the estimation of the counts-to-distance parameter α . Indeed, while the estimation of the α parameter is very stable on single chromosomes, we observe that this parameter α increases to -1 when the number of chromosomes increases for some non-metric methods. This has the effect of collapsing beads belonging to a chromosome together, while pushing beads belonging to other chromosomes away from one another.

To assess how well whole genome inference performs, we perform yet another stability experiment. We first infer a whole genome inference of a haploid *S. cerevisiae* dataset (Duan *et al.*, 2010). Then, we randomly subsample between 1 and 15 chromosomes and perform a whole genome 3D structure inference on this subsampled set of chromosomes. Finally, we assess the stability of the structure inference by computing the RMSD and Spearman correlation for each chromosome independently and taking the average to obtain a single RMSD and Spearman correlation score. Note that chromSDE and ShNeigh do not support multi-chromosome inference. We have thus excluded these two methods from these results. Results once again show that Pastis-NB is the most stable method in RMSD (Fig. 3C) and is a close second contender after Pastis-PM (Supplementary Fig. S5) for Spearman correlation. Supplementary Figure S7 shows the whole-genome reconstructions for the various methods.

4 Discussion and conclusion

We present in this work a new model, Pastis-NB, to infer a consensus 3D structure of the genome from Hi-C contact count data. We model interaction counts as negative binomial random variables, and we cast the inference as a likelihood maximization problem. Modeling counts as negative binomial random variables allows us to

better model the presence of overdispersion in Hi-C data, which we observed experimentally in Hi-C data from different organisms. Through extensive experiments on simulated and real Hi-C data, we showed that Pastis-NB consistently outperforms a representative set of competitive methods across a range of metrics. In particular, Pastis-NB yields remarkably stable and accurate results in the case of highly dispersed contact count data. This improvement is particularly striking at high resolution and at low coverage, with 3D models inferred much more robustly with Pastis-NB than with other methods. Yet, it is important to stress that it is not because a method is very stable that it is good or relevant. The striking differences in inferred 3D models between Pastis-NB and its closest contender, ShNeigh, highlight this limitation in measuring success with stability.

A limitation of Pastis-NB resides in the inference of consensus 3D models of chromosome architecture. Consensus models are not necessarily representative of the true folding of DNA in the cell. For example, a consensus model of *S. cerevisiae*'s genome will sometimes cluster telomeres together, while it is known that telomeres tether at the nuclear membrane. One should thus interpret these models with care. Yet consensus methods are powerful dimensionality reduction tools that can serve as an entry point for many analyses, such as data integration and visualization. In addition, such a consensus approach could be extended to single-cell Hi-C data, through appropriate modeling of the sparsity of the matrix via a zero-inflated negative binomial model and the addition of prior knowledge or biologically motivated constraints.

Funding

This work was supported by National Institutes of Health (NIH) [U54 DK107979 and UM1 HG011531]; and IRGA [G7H-IRG22B58].

Conflict of Interest: none declared.

References

- Anders, S. and Huber, W. (2010) Differential expression analysis for sequence count data. *Genome Biol.*, **11**, R106.
- Ay, F. *et al.* (2014) Three-dimensional modeling of the *P. falciparum* genome during the erythrocytic cycle reveals a strong connection between genome architecture and gene expression. *Genome Res.*, **24**, 974–988.
- Behr, J. *et al.* (2013) MITIE: simultaneous RNA-Seq-based transcript identification and quantification in multiple samples. *Bioinformatics*, **29**, 2529–2538.

- Ben-Elazar, S. et al. (2013) Spatial localization of co-regulated genes exceeds genomic gene clustering in the *Saccharomyces cerevisiae* genome. *Nucleic Acids Res.*, **41**, 2191–2201.
- Byrd, R.H. et al. (1995) A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.*, **16**, 1190–1208.
- Carty, M. et al. (2017) An integrated model for detecting significant chromatin interactions from high-resolution Hi-C data. *Nat. Commun.*, **8**, 15454.
- Cauer, A.G. et al. (2019). Inferring diploid 3D chromatin structures from Hi-C data. In: Huber, K. T. and Gusfield, D. (eds) *19th International Workshop on Algorithms in Bioinformatics (WABI 2019)*, Volume 143 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 11:1–11:13. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany.
- De, S. and Michor, F. (2011) DNA replication timing and long-range DNA interactions predict mutational landscapes of cancer genomes. *Nat. Biotechnol.*, **29**, 1103–1108.
- Deng, X. et al. (2015) Bipartite structure of the inactive mouse X chromosome. *Genome Biol.*, **16**, 152.
- Dixon, J.R. et al. (2012) Topological domains in mammalian genomes identified by analysis of chromatin interactions. *Nature*, **485**, 376–380.
- Duan, Z. et al. (2010) A three-dimensional model of the yeast genome. *Nature*, **465**, 363–367.
- Hu, M. et al. (2013) Bayesian inference of spatial organizations of chromosomes. *PLoS Comput. Biol.*, **9**, e1002893.
- Hu, M. et al. (2012) HiCNorm: removing biases in Hi-C data via Poisson regression. *Bioinformatics*, **28**, 3131–3133.
- Imakaev, M. et al. (2012) Iterative correction of Hi-C data reveals hallmarks of chromosome organization. *Nat. Methods*, **9**, 999–1003.
- Jin, F. et al. (2013) A high-resolution map of the three-dimensional chromatin interactome in human cells. *Nature*, **503**, 290–294.
- Kalhor, R. et al. (2012) Genome architectures revealed by tethered chromosome conformation capture and population-based modeling. *Nat. Biotechnol.*, **30**, 90–98.
- Kapilevich, V. et al. (2019) Chromatin 3D reconstruction from chromosomal contacts using a genetic algorithm. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, **16**, 1620–1626.
- Lesne, A. et al. (2014) 3D genome reconstruction from chromosomal contacts. *Nat. Methods*, **11**, 1141–1143.
- Lévy-Leduc, C. et al. (2014) Two-dimensional segmentation for analyzing Hi-C data. *Bioinformatics*, **30**, i386–i392.
- Li, F.-Z. et al. (2020) Chromatin 3D structure reconstruction with consideration of adjacency relationship among genomic loci. *BMC Bioinformatics*, **21**, 272.
- Li, J. et al. (2018) 3D genome reconstruction with ShRec3D+ and Hi-C data. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, **15**, 460–468.
- Lieberman-Aiden, E. et al. (2009) Comprehensive mapping of long-range interactions reveals folding principles of the human genome. *Science*, **326**, 289–293.
- Liu, D.C. and Nocedal, J. (1989) On the limited memory BFGS method for large scale optimization. *Math. Program*, **45**, 503–528.
- Nagalakshmi, U. et al. (2008) The transcriptional landscape of the yeast genome defined by RNA sequencing. *Science*, **320**, 1344–1349.
- Nikoloulopoulos, A.K. and Karlis, D. (2008) On modeling count data: a comparison of some well-known discrete distributions. *J. Stat. Comput. Simul.*, **78**, 437–457.
- Rao, S.P. et al. (2014) A 3D map of the human genome at kilobase resolution reveals principles of chromatin looping. *Cell*, **159**, 1665–1680.
- Rieber, L. and Mahony, S. (2017) miniMDS: 3D structural inference from high-resolution Hi-C data. *Bioinformatics*, **33**, i261–i266.
- Robinson, M.D. et al. (2010) edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, **26**, 139–140.
- Robinson, M.D. and Smyth, G.K. (2007) Moderated statistical tests for assessing differences in tag abundance. *Bioinformatics*, **23**, 2881–2887.
- Robinson, M.D. and Smyth, G.K. (2008) Small-sample estimation of negative binomial dispersion, with applications to SAGE data. *Biostatistics*, **9**, 321–332.
- Rousseau, M. et al. (2011) Three-dimensional modeling of chromatin structure from interaction frequency data using Markov chain Monte Carlo sampling. *BMC Bioinformatics*, **12**, 414.
- Ryba, T. et al. (2010) Evolutionarily conserved replication timing profiles predict long-range chromatin interactions and distinguish closely related cell types. *Genome Res.*, **20**, 761–770.
- Sexton, T. et al. (2012) Three-dimensional folding and functional organization principles of the *drosophila* genome. *Cell*, **148**, 458–472.
- Shen, Y. et al. (2012) A map of the *cis*-regulatory sequences in the mouse genome. *Nature*, **488**, 116–120.
- Tanizawa, H. et al. (2010) Mapping of long-range associations throughout the fission yeast genome reveals global genome organization linked to transcriptional regulation. *Nucleic Acids Res.*, **38**, 8164–8177.
- Tjong, H. et al. (2012) Physical tethering and volume exclusion determine higher-order genome organization in budding yeast. *Genome Res.*, **22**, 1295–1305.
- Varoquaux, N. (2019) Unfolding the genome: the case study of *P. falciparum*. *Int. J. Biostat.*, **15**.
- Varoquaux, N. et al. (2014) A statistical approach for inferring the 3D structure of the genome. *Bioinformatics*, **30**, i26–i33.
- Xekalaki, E. (2014) On the distribution theory of over-dispersion. *J. Stat. Distrib. Appl.*, **1**, 1–22.
- Yu, D. et al. (2013) Shrinkage estimation of dispersion in Negative Binomial models for RNA-seq experiments with small sample size. *Bioinformatics*, **29**, 1275–1282.
- Zhang, Y. et al. (2019) Large-scale 3D chromatin reconstruction from chromosomal contacts. *BMC Genomics*, **20** (Suppl 2), 186.
- Zhang, Y. et al. (2012) Spatial organization of the mouse genome and its role in recurrent chromosomal translocations. *Cell*, **148**, 908–921.
- Zhang, Z. et al. (2013). Inference of spatial organizations of chromosomes using semi-definite embedding approach and Hi-C data. In: *Proceedings of the 17th International Conference on Research in Computational Molecular Biology*, Volume 7821 of *Lecture Notes in Computer Science*, pp. 317–332, Springer, Berlin, Heidelberg.
- Zhu, G. et al. (2018) Reconstructing spatial organizations of chromosomes through manifold learning. *Nucleic Acids Res.*, **46**, e50.