Method

HiCRep: assessing the reproducibility of Hi-C data using a stratum-adjusted correlation coefficient

Tao Yang,¹ Feipeng Zhang,² Galip Gürkan Yardımcı,³ Fan Song,¹ Ross C. Hardison,^{1,4} William Stafford Noble,^{3,5} Feng Yue,^{1,6} and Qunhua Li^{1,2}

¹Bioinformatics and Genomics Program, Pennsylvania State University, University Park, Pennsylvania 16802, USA; ²Department of Statistics, Pennsylvania State University, University Park, Pennsylvania 16802, USA; ³Department of Genome Sciences, University of Washington, Seattle, Washington 98105, USA; ⁴Department of Biochemistry and Molecular Biology, Pennsylvania State University, University Park, Pennsylvania 16802, USA; ⁵Department of Computer Science and Engineering, University of Washington, Seattle, Washington 98105, USA; ⁶Department of Biochemistry and Molecular Biology, Pennsylvania State University College of Medicine, Hershey, Pennsylvania 17033, USA

Hi-C is a powerful technology for studying genome-wide chromatin interactions. However, current methods for assessing Hi-C data reproducibility can produce misleading results because they ignore spatial features in Hi-C data, such as domain structure and distance dependence. We present HiCRep, a framework for assessing the reproducibility of Hi-C data that systematically accounts for these features. In particular, we introduce a novel similarity measure, the stratum adjusted correlation coefficient (SCC), for quantifying the similarity between Hi-C interaction matrices. Not only does it provide a statistically sound and reliable evaluation of reproducibility, SCC can also be used to quantify differences between Hi-C contact matrices and to determine the optimal sequencing depth for a desired resolution. The measure consistently shows higher accuracy than existing approaches in distinguishing subtle differences in reproducibility and depicting interrelationships of cell lineages. The proposed measure is straightforward to interpret and easy to compute, making it well-suited for providing standardized, interpretable, automatable, and scalable quality control. The freely available R package HiCRep implements our approach.

[Supplemental material is available for this article.]

The three-dimensional (3D) genome organization across a wide range of length scales is important for proper cellular functions (Bickmore 2013; Dekker et al. 2013; Sexton and Cavalli 2015). At large distances, nonrandom hierarchical territories of chromosomes inside the cell nucleus are tightly linked with gene regulation (Misteli 2010). At a finer resolution, the interactions between distal regulatory elements and their target genes are essential for orchestrating correct gene expression across time and space (e.g., different tissues). A progression of high-throughput methods based on chromatin conformation capture (3C) (Dekker 2002) has emerged, including 4C (Simonis et al. 2006), 5C (Dostie et al. 2006), Hi-C (Lieberman-Aiden et al. 2009), ChIA-PET (Li et al. 2010), Capture Hi-C (Hughes et al. 2014; Mifsud et al. 2015), and HiChIP (Mumbach et al. 2016). These methods offer an unprecedented opportunity to study higher-order chromatin structure at various scales. Among them, the Hi-C technology and its variants are of particular interest due to their relatively unbiased genome-wide coverage and ability to measure chromatin interaction intensities between any two given genomic loci.

However, the analysis and interpretation of Hi-C data are still in their early stages. In particular, no sound statistical metric to evaluate the quality of Hi-C data has been developed. When biological replicates are not available, investigators often rely on either visual inspection of the Hi-C interaction heatmap or examination of the ratio of long-range interaction read pairs over the total sequenced reads (Dixon et al. 2012, 2015; Jin et al.

2013), but neither of these approaches is supported by robust statistics. When two or more biological replicates are available, it is a common practice to use either Pearson or Spearman correlation coefficients between the two Hi-C data matrices as a metric for data quality (Hu et al. 2012; Imakaev et al. 2012; Gorkin et al. 2014; Rao et al. 2014; Ay and Noble 2015; Dixon et al. 2015; Servant et al. 2015). However, Hi-C data have certain unique characteristics, including domain structures (such as topological association domain [TAD] and A/B compartments) and distance dependence, which refers to the fact that the chromatin interaction frequencies between two genomic loci, on average, decrease substantially as their genomic distance increases. Standard correlation approaches do not take into consideration these structures and may lead to incorrect conclusions. As we will demonstrate, two unrelated biological samples can have a high Pearson correlation coefficient, while two visually similar replicates can have a low Spearman correlation coefficient. It is also not uncommon to observe higher Pearson and Spearman correlations between unrelated samples than those between real biological replicates.

In this work, we developed HiCRep, a novel framework for assessing the reproducibility of Hi-C data that takes into account the unique spatial features of the data. HiCRep first minimizes the effect of noise and biases by smoothing the Hi-C matrix, and then it addresses the distance-dependence effect by stratifying Hi-C data according to their genomic distance. In particular, we developed

Corresponding authors: fyue@hmc.psu.edu, qunhua.li@psu.edu
Article published online before print. Article, supplemental material, and publication date are at <http://www.genome.org/cgi/doi/10.1101/gr.220640.117>.

© 2017 Yang et al. This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first six months after the full-issue publication date (see <http://genome.cshlp.org/site/misc/terms.xhtml>). After six months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

a stratum-adjusted correlation coefficient (SCC) as a similarity measure of Hi-C interaction matrices. SCC shares the similar range and interpretation as the standard correlation coefficients, making it easily interpretable. It can be used to assess the reproducibility of replicate samples or quantify the distance between Hi-C matrices from different cell types. Our framework also estimates confidence intervals for SCC, making it possible to infer statistical significance of the difference in reproducibility measurements. We applied our method to three different groups of publicly available Hi-C data sets to illustrate its power in distinguishing subtle differences between closely related cell lines and biological replicates and resolving interrelationship between different cell lineages or tissues.

Results

Spatial patterns in Hi-C data and their influence on reproducibility assessment

Unlike many other genomic data types, Hi-C data exhibit unique spatial patterns. One prominent pattern is the strong decay of interaction frequency as genomic distance increases between interaction loci, i.e., the distance dependence. This pattern is generally thought to result from nonspecific interactions, which are more likely to occur between loci at closer genomic distance than those at a greater distance (Fudenberg and Mirny 2012; Lajoie et al. 2015). It is found consistently in every Hi-C matrix and is one of the most dominant patterns in the matrix of interaction frequencies measured by Hi-C (Lajoie et al. 2015). This dependence on distance generates strong but spurious association between Hi-C matrices even when the samples are unrelated, as revealed by the high Pearson correlation between any two Hi-C matrices. As an example, we computed the Pearson correlations of Hi-C contact matrices between two biological replicates and between two unrelated cell lines, hESC and IMR90 (Dixon et al. 2012). Although samples from the same cell line are expected to be much more correlated to each other than samples from unrelated cell lines, the Pearson correlation shows little difference between samples from different cell types (a hESC sample and an IMR90 sample, $\rho = 0.92$) and biological replicates in hESC ($\rho = 0.91$) (Fig. 1A). Further investigation shows that the dependence pattern between the contact intensity and distance (Fig. 1B) is highly similar in hESC and IMR90, which creates the high, spurious correlation between the Hi-C samples from these two cell lines. Therefore, the Pearson correlation coefficient cannot distinguish real biological replicates from unrelated samples.

Another important pattern of Hi-C data is the domain structure in contact maps. These structures represent contiguous regions in which loci tend to interact more frequently with each

Figure 1. An illustration example. (A) Hi-C contact maps of the biological replicates of hESC and IMR90. (B) Relationship between genomic distance and the average contact frequency for the samples in A. Data are from Chromosome 22: 32000000–40000000.

other than with outside regions. Although the interactions within the structures can be highly variable between different cell types, the domain structures, such as topologically associating domains (TADs), are stable across cell types (Dixon et al. 2012; Rapkin et al. 2012; Nagano et al. 2013). Therefore, we expect a higher reproducibility at the domain level than at the individual contact level. This difference should be reflected in the reproducibility assessment. However, both Pearson and Spearman correlation coefficients only consider point interactions and do not take domain structures into account. A consequence of this is that Spearman correlation can be driven to low values by the stochastic variation in the point interactions and overlook the similarity in domain structures. As a result, two biological replicates that have highly similar domain structures may have a low Spearman correlation coefficient; conversely, a sample may have a higher Spearman correlation with an unrelated sample than with its biological replicates when the stochastic variation is high. For instance, despite the high similarity between the biological replicates in IMR90 and hESC, their Spearman correlations are only 0.47 and 0.37, respectively. However, the Spearman correlation between an IMR90 sample and a hESC sample (0.44) is higher than the correlation between the two hESC replicates (0.37), although there are many differences in the domain structures of the two cell lines. Therefore,

we need a more sophisticated evaluation metric to incorporate both structural aspects of variation for a better assessment of the reproducibility of Hi-C data.

Overview of the HiCRep method

We developed a novel two-stage approach to evaluate the reproducibility of Hi-C data (Fig. 2). The first stage is smoothing the raw contact matrix in order to reduce local noise in the contact map and to make domain structures more visible. The smoothing is accomplished by applying a 2D mean filter, which replaces the read count of each contact in the contact map with the average counts of all contacts in its neighborhood. In the second stage, we applied a stratification approach to account for the pronounced distance dependence in the Hi-C data. This stage proceeded in two steps. First, we stratified the smoothed chromatin interactions according to their genomic distance, and then we applied a novel stratum-adjusted correlation coefficient statistic (SCC) to assess the reproducibility of the Hi-C matrices. The SCC statistic was calculated by computing a Pearson correlation coefficient for each stratum (Fig. 2) and then aggregating the stratum-specific correlation coefficients using a weighted average, with the weights derived from the generalized Cochran–Mantel–Haenszel (CMH) statistic (Mantel 1963; Agresti 2012). The value of SCC ranges from -1 to 1 and can be interpreted in a way similar to the standard correlation. A great advantage of our approach is that we can derive the asymptotic variance of SCC and use it to assess statistical significance when comparing reproducibility from different samples. More detailed descriptions of the HiCRep method and the SCC statistic are presented in the Methods section.

Distinguishing pseudo, real, and nonreplicates

We first evaluated the performance of our method on samples whose expected levels of reproducibility are known: pseudoreplicates (PR), biological replicates (BR), and nonreplicates (NR). Biological replicates refer to two independent Hi-C experiments performed on the same cell types. Nonreplicates refer to Hi-C ex-

periments performed on different cell types. Pseudoreplicates are generated by pooling reads from biological replicates together and randomly partitioning them into two equal portions. The difference between two pseudoreplicates only reflects sampling variation, without biological or technical variation. Therefore, we expected the reproducibility of pseudoreplicates to be the highest, followed by biological replicates and then nonreplicates.

For testing, we first generated PR, BR, and NR using Hi-C data in the hESC and IMR90 cell lines (Methods; Dixon et al. 2012). We compared the performance of our method with Pearson correlation and Spearman correlation and investigated whether these metrics can distinguish PR, BR, and NR (Fig. 3A; Supplemental Table S1). For the hESC data set, our method correctly ranked the reproducibility of the three types of replicate pairs (PR>BR>NR), whereas Pearson and Spearman correlations both incorrectly ranked BR lower than one or more of the NRs. For the IMR90 data set, although all three methods infer the correct order of reproducibility, SCC separated BR from NR by a much larger margin (SCC: 0.19) than the other metrics (Pearson: 0.02; Spearman: 0.03).

The sequencing depths differ substantially for the hESC (replicate 1: 60M; replicate 2: 271M) and IMR90 (replicate 1: 201M; replicate 2: 153M) data sets. To ensure that these differences were not confounding our evaluations, we subsampled all the replicates to 60 million reads and repeated the same analysis. As shown in Figure 3A (blue dots), even with the same number of reads, Pearson and Spearman correlations still fail to distinguish real replicates from all nonreplicates. On the contrary, our method consistently ordered the reproducibility of replicates correctly, indicating that it can capture the intrinsic differences between the samples, even those that differ in sequencing depth.

We expanded this analysis to a larger Hi-C data set recently released by The ENCODE Project Consortium (2012). This data set consists of Hi-C data from eleven cancer cell lines, with two biological replicates for each cell type (Methods). For each cell type, we formed 20 nonreplicate pairs with the remaining 10 cell types and computed SCC, Pearson, and Spearman correlations for BR

and all NRs. As shown in Figure 3B and Supplemental Table S2, SCC clearly distinguishes BRs from NRs (a P -value = 1.665×10^{-15} , one-sided Kolmogorov–Smirnov [K-S] test), whereas the other two methods fail to do so (Pearson: P -value = 0.084; Spearman: P -value = 0.254, K-S test). Because the sequencing depth of the Hi-C data varies across cell types, we also examined the separation between BRs and NRs for each cell type. As shown in Figure 3C, SCC separates the BRs and NRs for all the cell types by a margin of at least 0.1, whereas the other two methods fail to separate them in more than half of the cell types (Supplemental Fig. S1). Furthermore, SCC illustrates a desirable correspondence to the sequencing depth. When the average sequencing depth between the biological replicates is relatively low (<30M), SCC monotonically increases with the sequencing depth; this behavior likely reflects insufficient coverage at the lower sequencing depths. In contrast, the value for SCC

Figure 2. A schematic representation of our method. (A) Step 1: smoothing; (B) Step 2: stratification.

Figure 3. Discrimination of pseudoreplicates (PR), biological replicates (BR), and nonreplicates (NR). (A) Reproducibility scores for the illustration example (hESC and IMR90 cell lines) in Figure 1. Red dots are the results in the original samples, and blue dots are the results after equalizing the sequencing depth in all samples. (B,C) Reproducibility scores for the BR and NR in the ENCODE 11 cancer cell lines. The triangle represents the score for a BR, and the box plot represents the distribution of the scores for NRs. (B) Reproducibility scores for BRs and NRs in all cell types. (C) SCC for BRs and the corresponding NRs in each cell type. From left to right, the cell lines are ordered according to the average sequencing depths of the biological replicates.

remains high and stable for greater sequencing depths (Fig. 3C), reflecting saturation of reproducibility and likely sufficient coverage. We investigate this property further in a subsequent section.

Evaluating biological relevance by constructing cell lineages

Next, we used HiCRep to infer the interrelationship between cell types on a cell lineage. Because this assessment requires the reproducibility measure to quantify the subtle differences between closely related cells, it serves as a biologically relevant approach to evaluating the accuracy of the reproducibility measure. More importantly, it also evaluates the potential of our method as a measure for quantifying the similarities or differences of Hi-C matrices in different cell or tissue types.

For this analysis, we used the Hi-C data in human embryonic stem (ES) cells and in four cell lineages derived from them (Dixon et al. 2015), namely, mesendoderm (ME), mesenchymal stem cells (MS), neural progenitor cells (NP), and trophoblast-like cells (TB), with two biological replicates for each cell type. Using the A/B compartments in Hi-C data, Dixon et al. (2015) inferred the distance to the parental ES cell from the nearest to the farthest as

ME, NP, TB, and MS (Fig. 4A, left). Importantly, the same relationships were also supported by our analysis of the previously published gene expression data (Xie et al. 2013) in the same cell types (Fig. 4A, right).

We first calculated the pairwise similarities between the 10 samples (two replicates in each cell type) using SCC, Pearson, and Spearman correlations (Supplemental Table S3). As shown in Supplemental Figure S2, SCC again provided the best separation between real replicates and nonreplicates among all three methods of comparison.

Next, we reconstructed the relationships among the cell lineages by performing hierarchical clustering based on the pairwise similarity scores. As shown in Figure 4B, the dendrogram constructed based on SCC precisely depicts the interrelationships: All the biological replicates are grouped together as terminal clusters, and the relationships between cell lines exactly follow the tree structure based on A/B compartments in Hi-C data (Dixon et al. 2015) and RNA-seq data in Xie et al. (2013) (Fig. 4A). The same results are obtained using the bias-corrected interaction matrices (Supplemental Fig. 3A). In contrast, the dendrograms constructed based on Pearson (Fig. 4C) and Spearman correlation coefficients (Fig. 4D) group several nonreplicates together and infer different relationships between some cell lines. For example, when using Pearson correlation, two ME replicates are not clustered together and NP is unexpectedly placed as the least related cell type to ES cells. When using Spearman correlation, an ES replicate is clustered with an ME replicate, and again, NP is unexpectedly predicted as the least related cell type to ES cells.

To further delineate how data smoothing contributes to HiCRep's performance, we also computed SCC on unsmoothed Hi-C matrices and Pearson and Spearman correlation coefficients on the smoothed Hi-C matrices obtained from our smoothing procedure. As shown in Supplemental Figure S3B, we observed that the SCC analysis on unsmoothed data no longer recapitulates the expected relationships among cell lineages, indicating that the smoothing stage is an indispensable component of HiCRep. Furthermore, we observed that our smoothing procedure improves the performance of Pearson- and Spearman-based approaches (Supplemental Fig. S3C,D), confirming its effectiveness. However, the improvement on Pearson- and Spearman-based approaches is not to the level achieved by SCC (Supplemental Fig. S3C,D). For example, the tree based on Pearson incorrectly groups the ES cell and ME cell replicates together, and the tree based on Spearman incorrectly places the NP cell closest to ES cell. In addition, the Pearson correlation based on the smoothed matrices shows very little difference (range = [0.96, 1]) across cell lines, making it hard to distinguish closely related samples. Together, this indicates that

Figure 4. Estimating interrelationship between the 10 samples in the human H1 ESC lineage. (A) The heatmap and lineage relationship between the ES cell and its five derived cells based on A/B compartments in Hi-C data (Dixon et al. 2015) and RNA-seq data in (Xie et al. 2013). (B–D) Estimated interrelationship based on the pairwise similarity score calculated using SCC (B), Pearson correlation (C), and Spearman correlation (D). Heatmaps show the similarity scores. Dendrograms resulted from a hierarchical clustering analysis based on the similarity scores. For easy visualization, the cell lines in the heatmaps are ordered according to their known distances to ES cells in A. A decreasing trend of scores is expected from left to right (from bottom to top, respectively) if the estimated interrelationship agrees with the known lineage.

neither SCC nor smoothing by itself can satisfactorily address the reproducibility issue, and both components are necessary for HiCRep to quantify biologically relevant differences between Hi-C contact maps.

We further expanded this analysis using the recently published Hi-C data in 14 human primary tissues and two cell lines (Supplemental Table S4; Schmitt et al. 2016). Because biological replicates are not available for all the samples, our analysis focused on quantifying the relationships between tissues or cells. Again, the lineage constructed based on SCC reasonably depicted the tissue and germ layer origins of the samples (Fig. 5A): hippocampus and cortex were grouped together; right ventricle and left ventricle were grouped together; and endodermal tissues such as pancreas, lung, and small bowel were placed in the same lineage. Neither Pearson nor Spearman correlation performed as well as SCC. For example, right and left ventricles were not grouped together by Spearman correlation (Fig. 5C). These results confirm the potential of our method as a measure for

quantifying the difference in Hi-C data between cell or tissue types.

HiCRep is robust to different choices of resolution

Depending on the sequencing depth, Hi-C data analysis may be performed at different resolutions. A good reproducibility measure should perform well despite the choice of resolution. To evaluate the robustness of our method, we repeated the clustering analysis for the human ES and ES-derived cell lineages using data processed at several different resolutions (i.e., 10, 25, 40, 100, 500 kb, and 1 Mb). Again, as shown in Figure 6 and Supplemental Table S5, SCC accurately inferred the expected relationships between ES and its derived cell lines at all resolutions considered, whereas Pearson and Spearman correlations inferred the expected relationships only at 500 kb and 1 Mb. Furthermore, unlike Pearson and Spearman correlations, whose values drastically change at different resolutions, the values of SCC remain in a consistent range

Figure 5. Estimated interrelationship for 14 human primary tissues and two cell lines in Schmitt et al. (2016). The dendrograms result from a hierarchical clustering analysis based on the pairwise similarity calculated using SCC (A), Pearson correlation (B), and Spearman correlation (C).

across all resolutions. The complete trees inferred by SCC at different resolutions (Supplemental Fig. S4) all agree with the expected relationship. These results confirm the robustness of our method to the choice of resolution.

Detecting differences in reproducibility due to sequencing depth

Sequencing depth is known to affect the signal-to-noise ratio and the reproducibility of Hi-C data (Lajoie et al. 2015). Insufficient coverage can reduce the reproducibility of a Hi-C experiment. As a quality control tool, a reproducibility measure is expected to be able to detect the differences in reproducibility due to sequencing depth. To evaluate the sensitivity of our method to sequencing depth, we downsampled all the samples in the H1 ES cell lineage (Dixon et al. 2015) to create a series of subsamples with different sequencing depths (25%, 50%, and 75% of the original sequencing depth). We then computed SCC for all subsamples. As shown in Figure 7A and Supplemental Table S6, SCC monotonically decreases with sequencing depth in all data sets. This confirms that our method can reflect the change of reproducibility between replicate experiments due to sequencing depth.

Furthermore, we investigate whether the reproducibility between Hi-C experiments inferred by SCC reflects the reproducibility at the level of significant contacts. To proceed, we identified significant contacts in these subsamples using Fit-Hi-C (Q -value cutoff = 0.05) (Ay et al. 2014). For each subsample, we computed the reproducibility at the level of significant contacts using the Jaccard index, i.e., the ratio of the number of shared significant contacts over the number of significant contacts identified in either replicate. As shown in Figure 7B and Supplemental Table S6, the Jaccard index monotonically increases with the SCC score in each cell line. In addition, the cell line (NP), which has a significantly lower SCC score than other cell lines, also shows a significantly lower level

of shared significant contacts. This confirms that our method can also reflect the change of reproducibility due to sequencing depth at the level of significant contacts.

Guiding the selection of the optimal sequencing depth

Having established that SCC can reflect the change of reproducibility due to the change of sequencing depth, we propose to use the saturation of SCC as a criterion to determine the most cost-effective sequencing depth that achieves a reasonable reproducibility. To illustrate how to use our method to determine the optimal sequencing depth, we created subsamples at a series of reduced sequencing depths from the Hi-C data in the H1 ES cell in Dixon et al. (2015) (original depth = 500 M) by downsampling. As shown in Figure 7C and Supplemental Table S7, SCC initially increases rapidly with the increase of sequencing depth when the number of total reads is less than 200 million (slope of line from 10% to 40% depth is 0.0591 per 100 M). It continues to increase at a reduced rate (slope of line from 40% to 70% depth is 0.01 per 100

Figure 6. Estimated similarity between the human H1 ES cell and its derived cells at different resolutions. (A) SCC; (B) Pearson correlation coefficient; and (C) Spearman correlation coefficient.

Figure 7. Detecting the change of reproducibility due to sequencing depth using SCC. (A) SCC of downsampled biological replicates (25%, 50%, 75%, and 100% of the original sequencing depth) for the five cell lines on the H1 ES cell lineage. (B) Relationship between SCC and Jaccard index, which measures the proportion of shared significant contacts identified by Fit-Hi-C between replicates for samples in A. (C) Saturation curves of SCC for data sets with different coverages. The SCC is plotted at different sub-samples (10%–90%) of the original samples with 90% confidence intervals. The blue dots represent H1 human ESC data (original sequencing depth = 500 M). The red dots represent the A549 data (original sequencing depth = 30 M).

M) and eventually reaches a plateau (slope of line from 70% to 100% depth is 0.002 per 100 M). To determine the lowest sequencing level that achieves similar reproducibility as the original data, we computed the 99% confidence intervals of SCC at each sequencing depths. Starting at 350 M (70% of the original depth), the confidence intervals of SCC between two adjacent depth levels overlap with each other, and the difference of SCC from that of the original depth is less than 0.004. This indicates that the reduced samples can achieve a similar level of reproducibility as the original one by using ~70% of the original depth for this data set. Further increase of sequencing depth beyond this point does not significantly improve reproducibility.

As a comparison, we performed a similar analysis using a data set with relatively low sequencing depth (30 M Hi-C reads from the A549 cell line). We observe that all the reduced samples have a significantly lower reproducibility than the original sample at the 99% significance level and show a steep increase of SCC throughout all downsample levels (Fig. 7C; Supplemental Table S7). From the 90% depth to the original depth, there is still an increase of SCC of 0.01, compared with 0.0015 for the hESC data set, suggesting that this data set may not reach saturation in reproducibility at its original sequencing depth. For this data set, further increase of sequencing depth may improve reproducibility.

Discussion

Although there has been a dramatic increase in the scope and complexity of Hi-C experiments, analytical tools for data quality control have been lacking. Current approaches for assessing Hi-C data reproducibility may lead to incorrect conclusions because they fail to take into consideration the unique spatial characteristics of Hi-C data. In this work, we developed a new method, HiCRep, for assessing the reproducibility of Hi-C contact frequency maps. By effectively taking account of the spatial features of Hi-C data, our reproducibility measure overcomes the limitations of Pearson and Spearman correlations and can differentiate the reproducibility of samples at a fine level. The empirical evaluation showed that our method distinguished subtle differences between closely related cell lines, biological replicates and pseudoreplicates, and it produced robust results at different resolutions.

The SCC statistic has several properties that make it well-suited as a reproducibility measure for providing standardized, interpretable, automatable, and scalable quality control. First, this statistic has a fixed scale of $[-1, 1]$, which makes it easy to standardize the quality-control process and compare reproducibility across samples. Second, it is intuitive and easy to interpret. It can be interpreted as a weighted average correlation coefficient over different interaction distances. This straightforward interpretation makes it accessible to

experimentalists. Third, our statistic is fast to compute and is directly applicable to the raw contact matrix. It is easily scalable for monitoring data quality for a large number of experiments. Furthermore, we also provide an estimator for the variance of this statistic, such that the statistical significance of the difference in reproducibility can be inferred. Using this estimator, we established a procedure to determine the sufficiency of sequencing depth.

In summary, we developed a novel method to accurately evaluate the reproducibility of Hi-C experiments. The presented method is a first step toward ensuring high reproducibility of Hi-C data. We also showed that this method can be used as a similarity measure for quantifying the differences in Hi-C data between different cell and tissue types. Thus, HiCRep is a valuable tool for the study of 3D genome organization.

Methods

Data sets

The data sets analyzed in this study were obtained from the public domain, as described below. Hi-C data sets used in this project can be visualized in the 3D genome browser (<http://3dgenome.org>).

We obtained the Hi-C data of human embryonic stem cells (hESCs) and human IMR90 fibroblasts from Dixon et al. (2012)

(GEO accession number: GSE35156). Each cell type has two biological replicates.

We obtained the Hi-C data of human embryonic stem (ES) cells and four human ES-cell-derived lineages, mesendoderm (ME), mesenchymal stem (MS) cells, neural progenitor (NP) cells and trophoblast-like (TB) cells from Dixon et al. (2015) (GEO accession number: GSE52457). Each cell type has two biological replicates.

We obtained the Hi-C data of 11 human cancer cell lines from the ENCODE data portal (Sloan et al. 2016) (https://www.encodeproject.org/matrix/?type=Experiment&status=released&assay_slims=3D+chromatin+structure&award.project=ENCODE&assay_title=Hi-C). This data set was produced by the Dekker laboratory. It includes cell lines of G401, A549, CAK12, PANC1, RPMI7951, T47D, NCIH460, SKMEL5, LNCaP, SKNMC, and SKNDZ. Each cell line has two biological replicates. The sequencing depths of the data sets can be found in Supplemental Table S8.

We obtained the Hi-C data of 14 human primary tissues from Schmitt et al. (2016) and Leung et al. (2015). The tissues include adrenal gland (GSM2322539), bladder (GSM2322540, GSM2322541), dorsolateral prefrontal cortex (GSM2322542), hippocampus (GSM2322543), lung (GSM2322544), ovary (GSM2322546), pancreas (GSM2322547), psoas muscle (GSM2322551), right ventricle (GSM2322554), small bowel (GSM2322555), spleen (GSM2322556), liver (GSM1419084), left ventricle (GSM1419085), and aorta (GSM1419086). The tissues were collected from four donors, each of which provides a subset of tissues. To minimize variation due to individual difference, we used the samples from the two donors with the largest number of tissues. If one tissue sample consists of multiple replicates from a single donor, the replicates were merged into a single data set. We obtained the GM12878 cell data from Selvaraj et al. (2013) (GSM1181867, GSM1181867) and the IMR90 cell data from Dixon et al. (2012) (GSM862724, GSM892307).

Data preprocessing

We generated the Hi-C contact matrices using the pipeline from Dixon et al. (2015). Briefly, the paired-end reads were first aligned to the hg19 reference genome assembly using BWA (Li and Durbin 2009). The unmapped reads were filtered, and potential PCR duplicates were removed using Picardtools (<https://broadinstitute.github.io/picard/>). We analyzed Hi-C reads mapped to human genome assembly hg19, because many of the published data sets were mapped to this assembly and this assembly has the deepest annotation of candidate functional noncoding sequences. Although the human genome assembly GRCh38 has improved contiguity in local regions, both assemblies are of high quality. Thus genome-wide analyses, such as our assessments of reproducibility, are not expected to differ significantly when the Hi-C reads are aligned to different assemblies. Importantly, our evaluation of reproducibility is not dependent on the exact locations to which reads map, but rather it uses the distances between two interacting sequences. Thus our metrics should be stable across similar assemblies.

For most analysis, we used 40-kb bins. To obtain contact maps at this resolution, we divided the genome into 40-kb bins as in Dixon et al. (2015) and obtained the interaction frequency by counting the number of reads falling into each pair of bins.

Our analysis only considered the intra-chromosomal interactions and only used the contacts within the range of 0–5 Mb in the reproducibility assessment. This range was chosen based on the observation that most of the A/B compartments have an interaction size of about 5 Mb, and interactions >5 Mb in distance are rare (<5% of reads) and highly stochastic. To evaluate the effect

of this parameter, we constructed the ES and its derived cell lineages using interactions in several ranges, including 0–4, 0–5, 0–6, 0–8, and 0–10 Mb (Fig. 4B, 0–5 Mb; Supplemental Fig. S5, others). The ranges of 0–4, 0–5, and 0–6 Mb gave the best results, confirming that 0–5 Mb is a reasonable choice. Only the bins with at least one count in at least one of the samples are kept for computing Pearson and Spearman correlations. All the data sets were preprocessed using the same procedure.

Our method is applicable to both raw and bias corrected data. However, to ensure the reproducibility assessment is free of assumptions made in the bias correction procedures (Cournac et al. 2012; Hu et al. 2012; Imakaev et al. 2012) and faithfully reflects the nature of the raw data, we chose to apply our method directly to raw data without bias correction. For the ES and its derived cell lines, we also applied our method to the bias-corrected matrices, in addition to the raw data, as a comparison. The bias correction was performed using the Iterative Correction and Eigenvector decomposition procedure (ICE) (Imakaev et al. 2012).

2D mean filter smoothing

Because the space of interactions surveyed by Hi-C experiments is very large, achieving sufficient coverage is still challenging. When samples are not sufficiently sequenced, the local variation introduced by under-sampling can make it difficult to capture large domain structures.

To handle this issue, we first smoothed the contact map before assessing reproducibility. Although smoothing will reduce the individual spatial resolution, it can improve the contiguity of the regions with elevated interaction, consequently enhancing the domain structures. It has been effective in commonly used Hi-C normalization methods (Yaffe and Tanay 2011; Imakaev et al. 2012).

We use a 2D mean filter to smooth the contact map. The filter replaces the read count of each contact in the contact map with the mean counts of all contacts in its genomic neighborhood. This filter is fast to compute and is effective for smoothing rectangular shapes (Davies 2012) like domain structures in Hi-C data. Specifically, let $C_{n \times n}$ denote a $n \times n$ contact map and c_{ij} denote the counts of the interaction between loci i and j . Given a span size $h > 0$, the smoothed contact map after passing an h th 2D mean filter is defined as follows:

$$x_{ij}(h) = \frac{\sum_{m=\max(1,i-h)}^{\min(i+h,n)} \sum_{l=\max(1,j-h)}^{\min(j+h,n)} C_{ml}}{(1+2h)^2}.$$

A visualization of the smoothing effect with different window sizes is shown in Supplemental Figure S6.

Selection of smoothing parameter

The span size h is a tuning parameter controlling the smoothing level. A very small h might not reduce enough local variation to enhance the boundaries of domain structures, whereas a large h will make the boundaries of domain structures blurry and limit the spatial resolution. Therefore, the optimal h should be adaptively chosen from the data.

To select h objectively, we developed a heuristic procedure to search for the optimal choice. Our procedure is designed based on the observation that the correlation between contact maps of replicate samples first increases with the level of smoothness and then plateaus when sufficient smoothness is reached. To proceed, we used a pair of reasonably deeply sequenced interaction maps as the training data. We randomly sampled 10% of the data 10 times. For each subsample, we computed the stratum-adjusted correlation coefficient (SCC,

described subsequently) at a series of h s in the ascending order and recorded the smallest h at which the increment of SCC was less than 0.01. The mode of h among the 10 subsamples was selected as the final span size. The detailed steps are shown in Algorithm 1 (Supplemental File S1).

Because the level of local variation in a contact map depends on the resolution used to process the data, the span size required to achieve sufficient smoothness varies according to resolution. Hence, a proper h for each resolution needs to be trained separately. However, at a given resolution, it is desirable to use the same h for all data sets, so that the downstream reproducibility assessment can be compared on the same basis. To reduce the chance of oversmoothing due to sparseness caused by insufficient coverage when training h , we used a deeply sequenced data set as training data.

Here, we obtained h in our analysis from the Human H1 ESC data set (Dixon et al. 2015). This data set was deeply sequenced (330 and 740 M reads for its two replicates) and had a reasonable quality (Dixon et al. 2015), making it suitable as training data. We processed the data using a series of resolutions (10, 25, 40, 100, 500 kb, and 1 Mb), and then we selected h for each resolution using the procedure described above. We obtained $h = 20, 11, 5, 3, 1$, and 0 for the resolution of 10, 25, 40, 100, 500 kb, and 1 Mb, respectively. These values were used throughout our study for all data sets at the corresponding resolutions. The robustness of our procedure was assessed using the Human H1 ESC data set and four derived cell lines (Results).

Stratification by distance

To take proper account of the distance effect in reproducibility assessment, we stratify the contacts by the genomic distance between their interaction loci. Specifically, let $X_{n \times n}$ be an $n \times n$ smoothed contact map at a resolution of bin size b . We compute the interaction distance for each contact x_{ij} as $d_{ij} = |j - i| \times b$ and then stratify the contacts by d_{ij} into K strata, $X_k = \{x_{ij} : (k-1)b < d_{ij} \leq kb\}$, $k = 1, \dots, K$. Here, we consider the interaction distance of 0–5 Mb. This leads to $K = 125$ for the bin size $b = 40$ kb. If x_{ij} is 0 in both samples, then it is excluded from the reproducibility assessment.

Stratum-adjusted correlation coefficient (SCC)

Our reproducibility statistic is motivated from the generalized Cochran–Mantel–Haenszel (CMH) statistic M^2 . The CMH statistic is a stratum-adjusted summary statistic for testing if two variables are associated while being stratified by the third variable (Agresti 2012), for example, the association between treatment and response stratified by age. Although originally developed for categorical data, it is also applicable to continuous data (Mantel 1963) and can detect consistent linear association across strata. However, the magnitude of M^2 depends on the sample size; therefore, it cannot be used directly as a measure of the strength of the association. When there is no stratification, the CMH statistic is related to the Pearson correlation coefficient ρ as $M^2 = \rho^2(N-1)$, where N is the number of observations (Agresti 2012). This relationship allows the strength of association summarized by M^2 to be represented using a measure that has a fixed scale and is comparable across different samples. However, ρ does not involve stratification. This motivates us to derive a stratum-adjusted correlation coefficient (SCC) to summarize the strength of association from the CMH statistic when there is stratification.

Derivation of stratum-adjusted correlation coefficient (SCC)

Let (X, Y) denote a pair of samples with N observations. The observations are stratified into K strata, and each stratum has N_k observations such that $\sum_{k=1}^K N_k = N$. Denote the observations in stratum k as $(x_{1k}, y_{1k}), \dots, (x_{N_kk}, y_{N_kk})$ and the corresponding random variables as (X_k, Y_k) , respectively. In our context, (x_{ik}, y_{ik}) are the smoothed counts of the i th contact on the k th stratum in the two contact maps X and Y . Let $T_k = \sum_{i=1}^{N_k} x_{ik}y_{ik}$, the CMH statistics is defined as

$$M^2 = \frac{\left[\sum_k [T_k - E(T_k)] \right]^2}{\sum_k \text{var}(T_k)}, \quad (1)$$

where $E(T_k)$ and $\text{var}(T_k)$ are the mean and variance of T_k under the hypothesis that X_k and Y_k are conditionally independent given the stratum,

$$E(T_k) = \frac{\sum_{i=1}^{N_k} x_{ik} \sum_{j=1}^{N_k} y_{jk}}{N_k}, \quad (2)$$

and

$$\begin{aligned} \text{var}(T_k) = & \frac{1}{N_k - 1} \left[\sum_{i=1}^{N_k} x_{ik}^2 - \frac{\left(\sum_{i=1}^{N_k} x_{ik} \right)^2}{N_k} \right] \\ & \times \left[\sum_{j=1}^{N_k} y_{jk}^2 - \frac{\left(\sum_{j=1}^{N_k} y_{jk} \right)^2}{N_k} \right]. \end{aligned} \quad (3)$$

To derive the stratum-adjusted correlation coefficient from the CMH statistic, write the Pearson correlation coefficient ρ_k for the k th stratum as $\rho_k = r_{1k}/r_{2k}$, where

$$r_{1k} = E(X_k Y_k) - E(X_k)E(Y_k) = \frac{\sum_{i=1}^{N_k} x_{ik}y_{ik}}{N_k} - \frac{\sum_{i=1}^{N_k} x_{ik} \sum_{j=1}^{N_k} y_{jk}}{N_k^2} \quad (4)$$

$$\begin{aligned} r_{2k} = & \sqrt{\text{var}(X_k)\text{var}(Y_k)} \\ = & \sqrt{\left[\frac{\sum_{i=1}^{N_k} x_{ik}^2}{N_k} - \left(\frac{\sum_{i=1}^{N_k} x_{ik}}{N_k} \right)^2 \right] \left[\frac{\sum_{j=1}^{N_k} y_{jk}^2}{N_k} - \left(\frac{\sum_{j=1}^{N_k} y_{jk}}{N_k} \right)^2 \right]}. \end{aligned} \quad (5)$$

It is easy to see that $r_{1k} = (T_k - ET_k)/N_k$ and $r_{2k} = \sqrt{(N_k - 1)\text{var}(T_k)}/N_k$. Then we can represent M^2 using ρ_k ,

$$M^2 = \frac{\left(\sum_{k=1}^K N_k r_{2k} \rho_k \right)^2}{\sum_{k=1}^K (N_k r_{2k})^2 / (N_k - 1)}. \quad (6)$$

Define

$$\rho_s = \frac{\sum_{k=1}^K N_k r_{2k} \rho_k}{\sum_{k=1}^K N_k r_{2k}} = \sum_{k=1}^K \left(\frac{N_k r_{2k}}{\sum_{k=1}^K N_k r_{2k}} \right) \rho_k, \quad (7)$$

then

$$M^2 = \frac{\rho_s^2}{\frac{1}{\left(\sum_{k=1}^K N_k r_{2k} \right)^2} \sum_{k=1}^K (N_k r_{2k})^2 / (N_k - 1)}. \quad (8)$$

Equation 8 shows that ρ_s^2 reflects the strength of association in M^2 . This strength relates to M^2 in a similar way as the Pearson correlation to M^2 in the case without stratification. As shown in Equation 7, ρ_s is a weighted average of the stratum-specific correlation coefficients, with weights $w_k = N_k r_{2k} / \sum_{k=1}^K N_k r_{2k}$ assigned according to the variance and sample size of a stratum. We call ρ_s the stratum-

adjusted correlation coefficient (SCC). Similar to standard correlation coefficients, it satisfies $-1 \leq \rho_s \leq 1$. A value of $\rho_s = 1$ corresponds to a perfect positive correlation, a value of $\rho_s = -1$ corresponds to a perfect negative correlation, and a value of $\rho_s = 0$ corresponds to no correlation.

The variance of ρ_s can be computed as

$$\text{var}(\rho_s) = \frac{\sum_{k=1}^K [(N_k r_{2k})^2 \text{var}(\rho_k)]}{\left(\sum_{k=1}^K N_k r_{2k}\right)^2}, \quad (9)$$

where $\text{var}(\rho_k)$ is the asymptotic variance for the Pearson correlation coefficient in a single stratum and can be computed using Fisher transformation (Fisher 1921) as follows: Let r_k be the sample version of ρ_k and Z_k be the Fisher transformation of r_k , i.e., $Z_k = (1/2) \ln[(1+r_k)/(1-r_k)]$, then Z_k is approximately normally distributed with mean $(1/2) \ln[(1+\rho_k)/(1-\rho_k)]$ and variance $1/(N_k - 3)$ (Fisher 1921). By the Delta method (Casella and Berger 2002), r_k is asymptotically normally distributed with mean ρ_k and the asymptotic variance as $\text{var}(\rho_k) = (1 - \rho_k^2)^2 / (N_k - 3)$.

The idea of obtaining an average correlation coefficient based on the CMH statistic has been explored in Rubenstein and Davis (1999) in the context of contingency tables with ordered categories. However, its derivation has several errors, which lead to a different statistic that ignores the sample size differences in different strata.

Variance stabilized weights

The downside for Equation 7 is that it is based on the implicit assumption in the CMH statistic that the dynamic ranges of X and Y are constant across strata. However, in Hi-C data, the read counts for contacts with short interaction distances have a much larger dynamic range than those with long interaction distances. As a result, the weights for the strata with large dynamic ranges will dominate Equation 7 owing to the large values of their r_{2k} . To normalize the dynamic range, we ranked the contact counts in each stratum separately and then normalized the ranks by the total number of observations N_k in each stratum, such that all strata shared a similar dynamic range. We then computed r_{2k} in the weights in Equations 7 and 9 using the normalized ranks, instead of the actual counts, i.e.,

$$r_{2k} = \sqrt{\text{var}\left[\frac{\text{Rank}(X_k)}{N_k}\right] \text{var}\left[\frac{\text{Rank}(Y_k)}{N_k}\right]} \quad (10)$$

The stratum-specific correlation ρ_k is still computed using actual values rather than ranks, as actual values have better sensitivity than ranks when there are a large number of low counts.

Software availability

We have implemented our method as an R package (R Core Team 2016). It is publicly available as the HiCRep package on GitHub (<https://github.com/qunhualilab/hicrep>) and in the [Supplemental Materials](#).

Acknowledgments

The authors thank Dekker's laboratory for providing the unpublished cancer cell line Hi-C data for our analysis. This work was supported by National Institutes of Health (NIH) training grant T32 GM102057 (CBIOS training program to The Pennsylvania State University), a Huck Graduate Research Innovation Grant (to T.Y.) and by NIH grants R01GM109453 (to Q.L.),

U01CA200060 (to F.Y.), R24DK106766 (to R.C.H. and F.Y.), U54HG006998 (to R.C.H.), and U41HG007000 (to W.S.N.).

Author contributions: Q.L. and F.Y. designed the study. Q.L. and T.Y. formulated the problem and developed the computational method and the SCC statistic. F.Z. provided assistance in the formulation of SCC statistic. Q.L., F.Y., and T.Y. designed the data analysis and performance evaluation. G.G.Y. and W.S.N. designed the data analysis for ENCODE cancer cell data. T.Y. implemented the method, developed the software package, and analyzed the data. T.Y. and F.S. optimized the software package. Q.L., T.Y., and F.Y. interpreted the results. R.C.H. provided assistance on the biological interpretation of the results. T.Y., R.C.H., F.Y., and Q.L. wrote the paper. All authors read and approved the final manuscript.

References

- Agresti A. 2012. *Categorical data analysis*. 3rd ed. Wiley, New York.
- Ay F, Noble WS. 2015. Analysis methods for studying the 3D architecture of the genome. *Genome Biol* **16**: 183.
- Ay F, Bailey TL, Noble WS. 2014. Statistical confidence estimation for Hi-C data reveals regulatory chromatin contacts. *Genome Res* **24**: 999–1011.
- Bickmore WA. 2013. The spatial organization of the human genome. *Annu Rev Genomics Hum Genet* **14**: 67–84.
- Casella G, Berger GL. 2002. *Statistical inference*, 2nd ed. Duxbury Press, Pacific Grove, CA.
- Cournac A, Marie-Nelly H, Marbouty M, Koszul R, Mozziconacci J. 2012. Normalization of a chromosomal contact map. *BMC Genomics* **13**: 436.
- Davies R. 2012. *Computer and machine vision: theory, algorithms, practicalities*. 4th ed. Academic Press, Oxford, UK.
- Dekker J. 2002. Capturing chromosome conformation. *Science* **295**: 1306–1311.
- Dekker J, Marti-Renom MA, Mirny LA. 2013. Exploring the three-dimensional organization of genomes: interpreting chromatin interaction data. *Nat Rev Genet* **14**: 390–403.
- Dixon JR, Selvaraj S, Yue F, Kim A, Li Y, Shen Y, Hu M, Liu JS, Ren B. 2012. Topological domains in mammalian genomes identified by analysis of chromatin interactions. *Nature* **485**: 376–380.
- Dixon JR, Jung I, Selvaraj S, Shen Y, Antosiewicz-Bourget JE, Lee AY, Ye Z, Kim A, Rajagopal N, Xie W, et al. 2015. Chromatin architecture reorganization during stem cell differentiation. *Nature* **518**: 331–336.
- Dostie J, Richmond TA, Arnaout RA, Selzer RR, Lee WL, Honan TA, Rubio ED, Krumm A, Lamb J, Nusbaum C, et al. 2006. Chromosome Conformation Capture Carbon Copy (SC): a massively parallel solution for mapping interactions between genomic elements. *Genome Res* **16**: 1299–1309.
- The ENCODE Project Consortium. 2012. An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**: 57–74.
- Fisher RA. 1921. On the probable error of a coefficient of correlation deduced from a small sample. *Metron* **1**: 3–32.
- Fudenberg G, Mirny LA. 2012. Higher-order chromatin structure: bridging physics and biology. *Curr Opin Genet Dev* **22**: 115–124.
- Gorkin DU, Leung D, Ren B. 2014. The 3D genome in transcriptional regulation and pluripotency. *Cell Stem Cell* **14**: 771–775.
- Hu M, Deng K, Selvaraj S, Qin Z, Ren B, Liu JS. 2012. HiCNorm: removing biases in Hi-C data via Poisson regression. *Bioinformatics* **28**: 3131–3133.
- Hughes JR, Roberts N, McGowan S, Hay D, Giannoulitou E, Lynch M, De Gobbi M, Taylor S, Gibbons R, Higgs DR. 2014. Analysis of hundreds of cis-regulatory landscapes at high resolution in a single, high-throughput experiment. *Nat Genet* **46**: 205–212.
- Imakaev M, Fudenberg G, McCord RP, Naumova N, Goloborodko A, Lajoie BR, Dekker J, Mirny LA. 2012. Iterative correction of Hi-C data reveals hallmarks of chromosome organization. *Nat Methods* **9**: 999–1003.
- Jin F, Li Y, Dixon JR, Selvaraj S, Ye Z, Lee AY, Yen CA, Schmitt AD, Espinoza CA, Ren B. 2013. A high-resolution map of the three-dimensional chromatin interactome in human cells. *Nature* **503**: 290–294.
- Lajoie BR, Dekker J, Kaplan N. 2015. The hitchhiker's guide to Hi-C analysis: practical guidelines. *Methods* **72**: 65–75.
- Leung D, Jung I, Rajagopal N, Schmitt A, Selvaraj S, Lee AY, Yen CA, Lin S, Lin Y, Qiu Y, et al. 2015. Integrative analysis of haplotype-resolved epigenomes across human tissues. *Nature* **518**: 350–354.
- Li H, Durbin R. 2009. Fast and accurate short read alignment with Burrows–Wheeler transform. *Bioinformatics* **25**: 1754–1760.
- Li G, Fullwood MJ, Xu H, Mulawadi FH, Velkov S, Vega V, Ariyaratne PN, Mohamed YB, Ooi HS, Tennakoon C, et al. 2010. ChIA-PET tool for

- comprehensive chromatin interaction analysis with paired-end tag sequencing. *Genome Biol* **11**: R22.
- Lieberman-Aiden E, van Berkum NL, Williams L, Imakaev M, Ragoczy T, Telling A, Amit I, Lajoie BR, Sabo PJ, Dorschner MO, et al. 2009. Comprehensive mapping of long-range interactions reveals folding principles of the human genome. *Science* **326**: 289–293.
- Mantel N. 1963. Chi-square tests with one degree of freedom; extensions of the Mantel-Haenszel procedure. *J Am Stat Assoc* **58**: 690–700.
- Mifsud B, Tavares-Cadete F, Young AN, Sugar R, Schoenfelder S, Ferreira L, Wingett SW, Andrews S, Grey W, Ewels PA, et al. 2015. Mapping long-range promoter contacts in human cells with high-resolution capture Hi-C. *Nat Genet* **47**: 598–606.
- Misteli T. 2010. Higher-order genome organization in human disease. *Cold Spring Harb Perspect Biol* **2**: a000794.
- Mumbach MR, Rubin AJ, Flynn RA, Dai C, Khavari PA, Greenleaf WJ, Chang HY. 2016. HiChIP: efficient and sensitive analysis of protein-directed genome architecture. *Nat Methods* **13**: 919–922.
- Nagano T, Lubling Y, Stevens TJ, Schoenfelder S, Yaffe E, Dean W, Laue ED, Tanay A, Fraser P. 2013. Single-cell Hi-C reveals cell-to-cell variability in chromosome structure. *Nature* **502**: 59–64.
- R Core Team. 2016. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>.
- Rao SS, Huntley MH, Durand NC, Stamenova EK, Bochkov ID, Robinson JT, Sanborn AL, Machol I, Omer AD, Lander ES, et al. 2014. A 3D map of the human genome at kilobase resolution reveals principles of chromatin looping. *Cell* **159**: 1665–1680.
- Rapkin LM, Anchel DR, Li R, Bazett-Jones DP. 2012. A view of the chromatin landscape. *Micron* **43**: 150–158.
- Rubenstein LM, Davis CS. 1999. Estimation of the average correlation coefficient for stratified bivariate data. *Stat Med* **18**: 567–580.
- Schmitt AD, Hu M, Jung I, Xu Z, Qiu Y, Tan CL, Li Y, Lin S, Lin Y, Barr CL, et al. 2016. A compendium of chromatin contact maps reveals spatially active regions in the human genome. *Cell Rep* **17**: 2042–2059.
- Selvaraj S, R Dixon J, Bansal V, Ren B. 2013. Whole-genome haplotype reconstruction using proximity-ligation and shotgun sequencing. *Nat Biotechnol* **31**: 1111–1118.
- Servant N, Varoquaux N, Lajoie BR, Viara E, Chen CJ, Vert JP, Heard E, Dekker J, Barillot E. 2015. HiC-Pro: an optimized and flexible pipeline for Hi-C data processing. *Genome Biol* **16**: 259.
- Sexton T, Cavalli G. 2015. The role of chromosome domains in shaping the functional genome. *Cell* **160**: 1049–1059.
- Simonis M, Klous P, Splinter E, Moshkin Y, Willemsen R, de Wit E, van Steensel B, de Laat W. 2006. Nuclear organization of active and inactive chromatin domains uncovered by chromosome conformation capture-on-chip (4C). *Nat Genet* **38**: 1348–1354.
- Sloan CA, Chan ET, Davidson JM, Malladi VS, Strattan JS, Hitz BC, Gabdank I, Narayanan AK, Ho M, Lee BT, et al. 2016. ENCODE data at the ENCODE portal. *Nucleic Acids Res* **44**: D726–D732.
- Xie W, Schultz MD, Lister R, Hou Z, Rajagopal N, Ray P, Whitaker JW, Tian S, Hawkins RD, Leung D, et al. 2013. Epigenomic analysis of multilineage differentiation of human embryonic stem cells. *Cell* **153**: 1134–1148.
- Yaffe E, Tanay A. 2011. Probabilistic modeling of Hi-C contact maps eliminates systematic biases to characterize global chromosomal architecture. *Nat Genet* **43**: 1059–1065.

Received January 22, 2017; accepted in revised form August 7, 2017.

HiCRep: assessing the reproducibility of Hi-C data using a stratum-adjusted correlation coefficient

Tao Yang, Feipeng Zhang, Galip Gürkan Yardimci, et al.

Genome Res. 2017 27: 1939-1949 originally published online August 30, 2017

Access the most recent version at doi:[10.1101/gr.220640.117](https://doi.org/10.1101/gr.220640.117)

Supplemental Material

<http://genome.cshlp.org/content/suppl/2017/10/06/gr.220640.117.DC1>

References

This article cites 38 articles, 5 of which can be accessed free at:
<http://genome.cshlp.org/content/27/11/1939.full.html#ref-list-1>

Creative Commons License

This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first six months after the full-issue publication date (see <http://genome.cshlp.org/site/misc/terms.xhtml>). After six months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

Email Alerting Service

Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#).

Affordable, Accurate
Sequencing.

To subscribe to *Genome Research* go to:
<https://genome.cshlp.org/subscriptions>