Journal Title Here, 2022, pp. 1–9

doi: DOI HERE
Advance Access Publication Date: Day Month Year
Paper

PAPER

L_0 segmentation enables data-driven concise representations of diverse epigenomic data

Ali Tuğrul Balcı[1,2] and Maria Chikina[1,*]

[1]Computational and Systems Biology, University of Pittsburgh and [2]Joint Carnegie Mellon-University of Pittsburgh Ph.D. Program in Computational Biology
[*]Corresponding author: mchikina@pitt.edu

Abstract

Motivation: Epigenetic assays using next-generation sequencing (NGS) have furthered our understanding of the functional genomic regions and the mechanisms of gene regulation. However, a single assay produces billions of data represented by nucleotide resolution signal tracks. The signal strength at a given nucleotide is subject to numerous sources of technical and biological noise and thus conveys limited information about the underlying biological state. In order to draw biological conclusions, data is typically summarized into higher order patterns. Numerous specialized algorithms for summarizing epigenetic signal have been proposed and include methods for peak calling or finding differentially methylated regions. A key unifying principle underlying these approaches is that they all leverage the strong prior that signal must be locally consistent. **Results:** We propose L_0 segmentation as a universal framework for extracting locally coherent signals for diverse epigenetic sources. L_0 serves to both compress and smooth the input signal by approximating it as piece-wise constant. We implement a highly scalable L_0 segmentation with additional loss functions designed for NGS epigenetic data types including Poisson loss for single tracks and binomial loss for methylation/coverage data. We show that the L_0 segmentation approach retains the salient features of the data over a wide range of compression values and can identify subtle features, such as transcription end sites, missed by other analytic approaches. **Availability:** Our approach is implemented as an R package "l01segmentation" with a C++ backend. Available at https://github.com/boooooogey/l01segmentation.

Key words: epigenetics, optimization, compression

Introduction

Using next-generation sequencing (NGS) to profile the epigenetic states of biological systems has rapidly gained popularity. Large scale efforts to profile the epigenome, such as ENCODE [1] and Roadmap [2], have paved the way for these assays being used routinely to investigate diverse biological problems from cancer biology to immune cell development.

Once the data is aligned, epigenetic assays can be represented as (possibly normalized) read densities along the chromosome. Without further processing, each experiment contains on the order of 3×10^9 values but many downstream analyses make use of representations with a much smaller footprint. With a few important exceptions, epigenetic assays produce signals that have peak-like structure, that is, sequencing reads pile up in specific chromosomal locations representing a small fraction of the genome. Such assays can be succinctly represented by recording the position and value of the peak. The MACS [3] peak discovery software is widely used by consortium projects. Importantly, some histone marks, such as K27me3 and K27me36, do not conform to the peak-like structure but instead cover broad regions at the scale of genes. In these cases the approach taken by MACS is to heuristically merge narrow peaks [4] forming so-called "broad peaks".

Sequence-based methylation is unique among epigenetic assays in that it reads out a covalent modification of the DNA and the output data does not have not have natural "peaks" that are the hallmark of other epigenetic tracks but is rather dispersed through the genome.

Despite not having the peak structure the methylation levels at adjacent sites show a high level of dependence, as they are likely to be governed by a common regulatory process, such as the binding of a TF or polymerase or gene transcription. For experimental designs that are focused on specific biological contrasts, this local coherence of methylation is exploited by combining differentially methylated sites (DMSs) into stretches of differentially methylated regions (DMRs) [5, 6, 7, 8, 9].

Altogether, the key property of epigenetic data is that while the raw data is reported at nucleotide resolution nearby chromosomal positions are likely to be produced by the same biological or experimental process and this fact can be exploited to define narrow peaks, broad peaks, or DMRs. While these analysis approaches are popular and effective, they rely on specific assumptions, tunable parameters, or, in the case of

© The Author 2022. Published by Oxford University Press. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

	Implementation Availability	Probability Distribution	Data Format Supported	Method class
CROCS [10]	Yes	Gaussian, Poisson, Negative Binomial	ChIP-seq	Breakpoint detection with structured constraints (originally proposed as gfpop[11]).
PELT [12]	Yes	Flexible	Flexible	Break-point detection with pruned dynamical programming
Yokoyama et al [13]	No	Binomial	WGBS	Application of PELT
Herrmann et al [14]	No	Binomial (signal vs. input)	ChIP-seq	Bayesian MCMC
Chen et al [15]	No	Gaussian	ChIP-seq	Bayesian MCMC
Xing et al [16]	No	Poisson	ChIP-seq	Hidden-markov model

Table 1. Comparison of existing segmentation methods.

DMRs, on knowledge of experimental design. In this work we propose a general-purpose framework that uses efficient L_0 segmentation and assay specific loss functions to represent any NGS epigentic experiment as piece-wise constant while making no a priori assumption regarding the expected patterns in the data.

Our Approach

A standard way of enforcing that signals are locally consistent is by penalizing the difference in adjacent values. Given an input vector y_i we seek to approximate it with a locally smoothed vector $\boldsymbol{\beta}$ by solving an optimization problem of the following general type:

$$\boldsymbol{\beta} = \underset{\beta \in \mathbb{R}^n}{\operatorname{argmax}} \sum_{i=1}^{N} e(\beta_i, y_i) - \lambda \sum_{i=1}^{N-1} \Omega\left(\beta_i - \beta_{i+1}\right) \qquad (1)$$

Here $e(\beta, y)$ is a reconstruction error term, and $\Omega(\beta_i - \beta_{i+1})$ is a regularization term that penalizes the differences between consecutive values.

If we take squared error and the penalty to be the L_1 norm this problem is known as fused-Lasso. This penalty drives the difference between successive elements of β to 0 and hence the approximation solution is piece-wise constant. Because of this property, the fused-Lasso can also be seen as a change-point detection algorithm.

In this work we consider a different penalty function, the L_0 error. Instead of penalizing the sum of the absolute values of the differences, this simply penalizes the total number of change points regardless of their values. The L_0 formulation has important performance and algorithmic consequences. Because all change points are penalized equally, the L_0 formulation can produce much sharper breaks, which, as we show, allows for more efficient segmented representations. Another attractive property of the L_0 formulation is that for the optimal solution, the value of the fused segment is purely a function of the individual data points in the segment (in the case of squared and Poisson error, it is the mean) and does not depend on the adjacent segments.

In terms of algorithmic considerations, unlike the L_1 formulation, the L_0 formulation is not convex. However, recent algorithmic advances allow for the solution to be found in empirical linear time [17].

The algorithm proposed by Johnson [17] has theoretical $\mathcal{O}\left(n^2\right)$ time and space complexity. However, in practice, the execution does not reach this theoretical limit and instead scales linearly in both time and space with the length of the input signal. The empirical linear space and time usage of the algorithm make the application to large epigenetic signals feasible.

Our implementation, which extends the approach described for solving L_0 segmentation [17], uses error functions derived from different probability distributions. We implement the standard Gaussian distribution (squared error) as well as distributions that more accurately model epigenetic count data, such as Poisson for single-track data and binomial for double-track data counts (i. e. methylation and coverage). We find empirically that the L_0 segmentation scales linearly when applied to biological data. In Fig. 1, we plot the time it takes to solve the L_0 segmentation problem with different probability distributions for various input sizes.

For completeness our implementation also includes functions that solve the L_1-fused approximation problem, another state-of-the-art ultra-fast algorithm which was also described in [17]. However, we will demonstrate the the L_0 formulation is superior in detecting relevant biological signals.

The empirical performance of our implementation both scales linearly and is overall very efficient, taking just 0.1 and 1 second (for Poisson and Binomial respectively) on 10^6 input. The scalability makes it possible to efficiently compress epigentic data genome-wide. For example, reading in the entire nucleotide-level Chip-seq signal of Pol II from the BigWig file, collating it into 20 base pair bins, applying L_0 segmentation, and storing the segmentation in a BedGraph file takes ˜90 seconds with 4 cores. Short-range local binning is reasonable for ChIP-seq data as the signal is not inherently nucleotide resolution. On the other hand, methylation data is single nucleotide, and binomial error is more computationally demanding. However, this is mediated by the fact that most methylation analysis is done on CpGs, and there are only ˜30 million CpGs in the human genome.

Relationship with prior work

A number of unsupervised approaches for segmenting epigenetic data have been proposed. These cover a diverse range of modeling assumptions and algorithmic approaches. Our work builds on these efforts in several important ways. Some previous methods have relied on Bayesian models using MCMC for inference. While this approach allows for modeling flexibility, it is too time-consuming to be applied genome-wide. Our approach, on the other hand, is based on penalizing breakpoints and is conceptually similar to PELT and CROCS. However, the algorithmic approach is distinct, resulting in a significant speedup (see Fig. 1). Moreover, we implement several useful error functions and can segment both single-track and double-track (e.g. methylation) data. We also provide a simple R package interface which allows direct segmentation of BigWig files.

We also note that in the context of methylation data analysis, it is common to use segmentation-type approaches to combine differentially methylated sites into differentially methylated regions [6, 18, 7, 8, 9, 19]. These approaches rely on

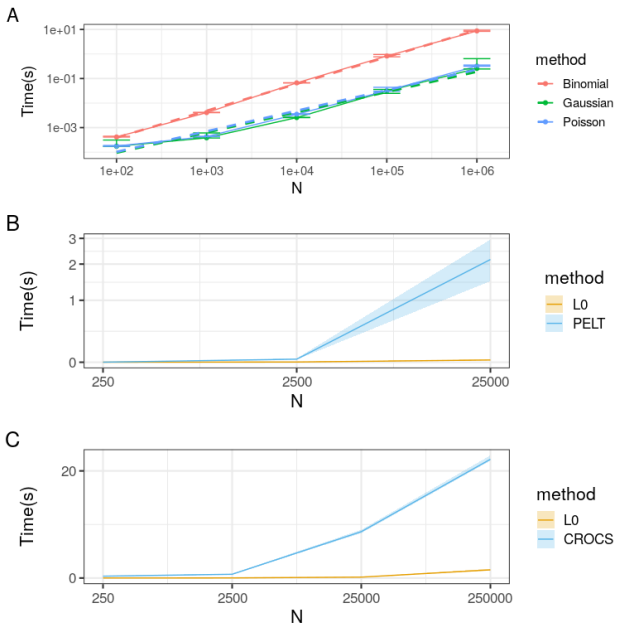

Fig. 1. A: Binomial, Gaussian, and Poisson L_0 segmentation execution times are plotted in log-space (both x and y) with respect to various segment lengths. The data is simulated using the corresponding probability distributions. Even though the theoretical worst case complexity of the algorithm is $O(n^2)$, in practice, the execution time scales linearly with the length of the input. **B** and **C:** comparison of the running time of our proposed method against similar approaches. CROCS is used for epigenetic data. PELT is the algorithm used in the Yokoyama et al. approach for segmenting methylation data

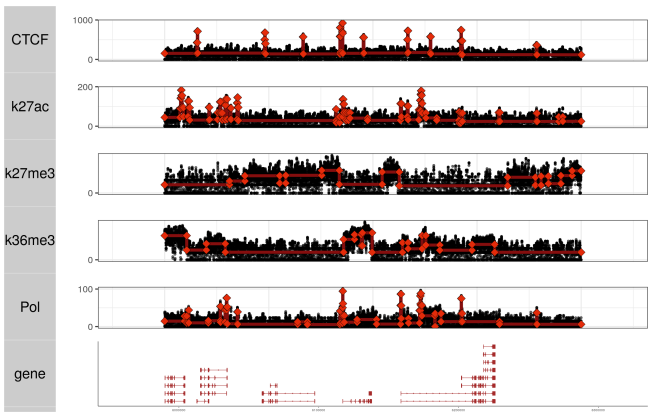

Fig. 2. Poisson L_0 segmentation on various epigenetic tracks is shown. The hyperparameter (λ) is chosen automatically using cross validation.

test statistics generated at the site level and generally cannot be applied out-of-the-box to the unsupervised segmentation of genome-wide signal.

L_0 segmentation retains salient genomic features

To illustrate features of L_0 segmentation, we plot segmentation results across tracks with diverse features that represent different signal classes: CTCF and K27ac for narrow peaks, K27me3 and K36me for broad peaks, and Pol II, which shows a complex signal pattern composed of high narrow peaks and low broad peaks. Each track was segmented with L_0 Poisson loss with the regularization parameter determined by optimizing "offset" cross-validation, a specialized cross-validation procedure that accounts for nuisance autocorrelation that arises from ChIP-seq fragments (see Methods).

We find that while neither the L_0 segmentation nor the cross-validation procedure is optimized for a particular signal type, our approach automatically fits the expected patterns to each data type, effectively finding narrow peaks and broad peaks.

To demonstrate these qualitative observations in a quantitative evaluation framework, we make use of raw and processed datasets from the NIH Roadmap Epigenomics Consortium. The database provides an organized compilation of high quality uniformly processed raw epigenetic data as well as output form standard analysis tools such as MACS or ChromHMM [3, 20]. Using these well-established analysis tools as a standard give us the opportunity to evaluate if we can distill different raw epigenetic tracks into a concise summary using a single principled algorithm. We focus on two cell lines that had all the relevant data types available: fibroblast cells (IMR90) and embryonic stem cells (H1).

The L_0 segmentation provides a reduced representation of a signal under the assumption that it is locally constant. As such, it doesn't explicitly optimize specific features. Alternative approaches to compressing epigenetic datasets are equal-width local binning or using specialized algorithms to explicitly detect specific features of interest such as "peaks". The MACS software [3] is highly optimized for peak discovery, and we use the MACS peak calls as a silver standard in order to evaluate the performance of L_0 compressed representations, which are in fact signal shape agnostic.

Specifically, we evaluate if reduced representations capture the peak-like features with high fidelity by formalizing a notion of representation efficiency, that is the trade of between representation size and signal capture which for segmentation methods can be controlled by varying λ. Plotting the signal fidelity (which we define in multiple ways) against representation size for different compressed representations allows us to measure their efficiency independent of specific hyperparameter choices. For our peak-based analysis, we evaluate L_0 segmentation (our proposed solution), the convex L_1 segmentation, and equal-size binning (a baseline that is commonly used to reduce epigenetic tracks in a data-agnostic way). We report the representation size in terms of "compression ratio," which is the input length divided by the number of constant segments in the output.

In order to quantify how well the reduced representations retain peak features (signal fidelity), we formulate two different evaluations. The first one is based on computing the median fold change under the MACS peak regions after data compression. At the extreme value of compression ratio =1 (no compression) this just corresponds to the median fold change for the raw data. As the compression ratio is increased this number will eventually decrease as peak regions are merged with adjacent regions. The results of this evaluation are plotted in Fig. 3 (A,B). We find that according to this metric the L_0 method is able to maintain accuracy up to 10,000-fold compression while both the binning and L1 show considerable drop-offs in performance much earlier.

The second evaluation is based on quantifying to what extent the reduced representation approximated the boundaries (as opposed to the values) of the peak calls. For each MACS peak we get the maximum Jaccard index across all possible overlapping segment region. The distributions are plotted in

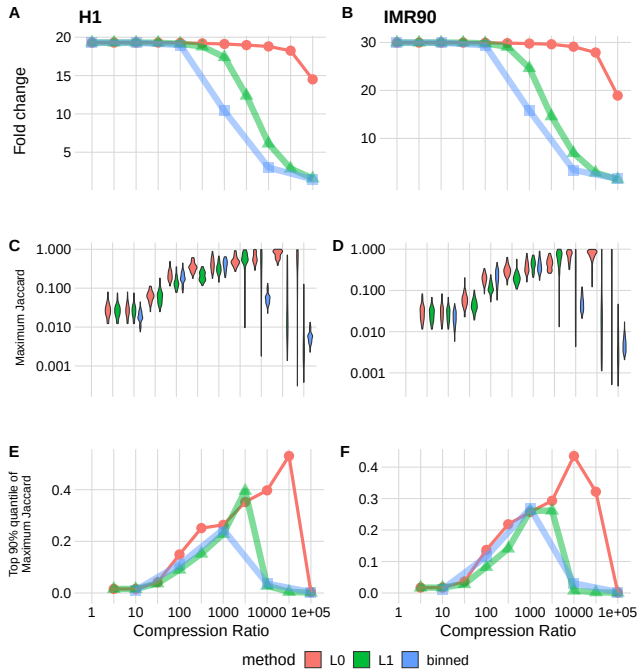

Fig. 3. Comparison of L_0Poisson (red), L1 Poisson (green), and fixed size binning (blue) with respect to the various compression ratios on the DNAse track from the H1 and IMR90 cell lines. The methods are applied on a 10M base-pair segment extracted from each cell line's Chromosome 1. We evaluated the methods using 937 peaks for H1 and 1400 peaks for IMR90 discovered by MACS within this segment. **A and B:** The ratio of the mean signal within peak regions to the mean background signal after segmentation. **C and D:** Maximum Jaccard Index distributions for the methods at different compression ratios. Jaccard indices are calculated between peaks and the segments discovered by the methods. **E and F:** Median of the distributions shown in C and D.

Fig. 3 C,D and 90% quantile trends are plotted in the third row in panels E and F. In this evaluation, we see that the input (compression ratio =1), where each point corresponds to a segment is actually not optimal . This is expected as for the input each point is a single nucleotide and thus will not have a high Jaccard index with peak regions that are on the order of hundreds of nucleotides long. Instead, the optimal is reached at an intermediate value for all methods though with widely different compression ratios. Considering the H1 cell-line (C,E), local binning achieves optimal signal at a compression ratio of about 1000-fold, where 90% of peaks have a Jaccard concordance value of 0.2. For L_1 the corresponding values are about 3000-fold and 0.4, a considerable improvement. On the other hand, L_0 achieves both the best quantile value of 0.5 and does so at a substantially higher compression rate of 3×10^4.

We also compare the compressed representation quality for three different epigenetic tracks (DNAse, H3K27ac, and H3K27me3) using the three different methods with respect to enrichment from ChromHMM states at a compression ratio of 10^4, as shown in Fig. 4. Importantly, we include H3K27me3, which is a repressive mark that does not have peak structure but rather covers long genomics regions on the scale of genes. For the evaluation, the ChromHMM regions provided by the Roadmap project are fixed, and we evaluate the enrichment of signal within a specific ChromHMM class using either the raw input (as is done in the original ChromHMM paper) or the compressed representation. We find that across the three compression methods, L_0 retains ChromHMM enrichment that is nearly identical to the input data (see Fig. 4 B).

Finally, we focus on segmenting Pol II tracks, which have complex signal shapes that cannot be easily captured in a peak-finding framework. As can be seen in the example shown in Fig. 5 B the segmentation procedure identifies both the sharp peak at the transcription start site and the broad transcript associated regions. In particular, we expect that a correct segmentation of this signal should place a breakpoint near the transcription. In Fig. 5 A, we quantify this via the median distance (MD) between a TES and the closest breakpoint using a variety of different data representations. We consider segmented data using either L_0 or L_1 penalty and Poisson or Gaussian distributions. The degree of compression is adjusted by the penalty parameter. We also consider the MACS peak based representation and vary the total number of peak based segments by adjusting the FDR cutoff. We demonstrate that the L0 Poisson model consistently achieves the lowest MD, with L0 Gaussian having the next best performance. In Fig. 5 B, we provide some examples of Pol II tracks and corresponding segmentations, demonstrating that L_0 segmentation effectively finds TES-associated breakpoints which are missed by the peak-based analysis (green) .

In summary, we demonstrate that the L_0 Poisson compression is highly effective in retaining the salient signals from diverse signal track types and in a variety of different evaluations. Using L_0 compression, it is possible to reduce the complexity of the epigenetic signal by 10,000-fold while retaining features important for downstream interpretation. Importantly, this approach is capable of capturing complex signal shapes such as the ones produced by Pol II with high biological fidelity. Thus, L_0 compression is a viable hypothesis-free alternative for learning meaningful reduced representations of epigenetic data.

L0 compression for methylation data

DNA methylation is a widely used assay type to profile the epigenetic state of cells and tissues. DNA methylation is unique among epigenetic signals in terms of biology, assay techniques, and analysis. Methylation is a covalent modification of the DNA which can be measured directly with bisulfite conversion and sequencing and does not require complex biochemical steps needed to perform chromatin-based assays. Bisulfite conversion can be used to add a methylation readout to other sequencing-based assays, as is done with single-cell HiC [21].

Assays to measure methylation in high throughput via probe hybridization have been in wide use for over a decade and have been applied to diverse biological problems, from cancer characterization to epidemiological studies [22, 23]. The wide adoption of array based technologies has revealed how DNA methylation reports not just the state but also the developmental and environmental history of the cell and organism, as exemplified by the various methylation clock models (see [24] for review).

Because methylation is a covalent modification, it is relatively stable and simple to assay. Methylation models are reproducible across cohorts, tissues, and assay platforms [24, 25, 26] and can be constructed for large numbers of phenotypes [27]. These properties of methylation make it an attractive front-line epigenetic assay.

Unlike other epigenetic tracks, methylation is truly nucleotide resolution as it directly measures the methylation status of individual cytosines. Methylation data is reported

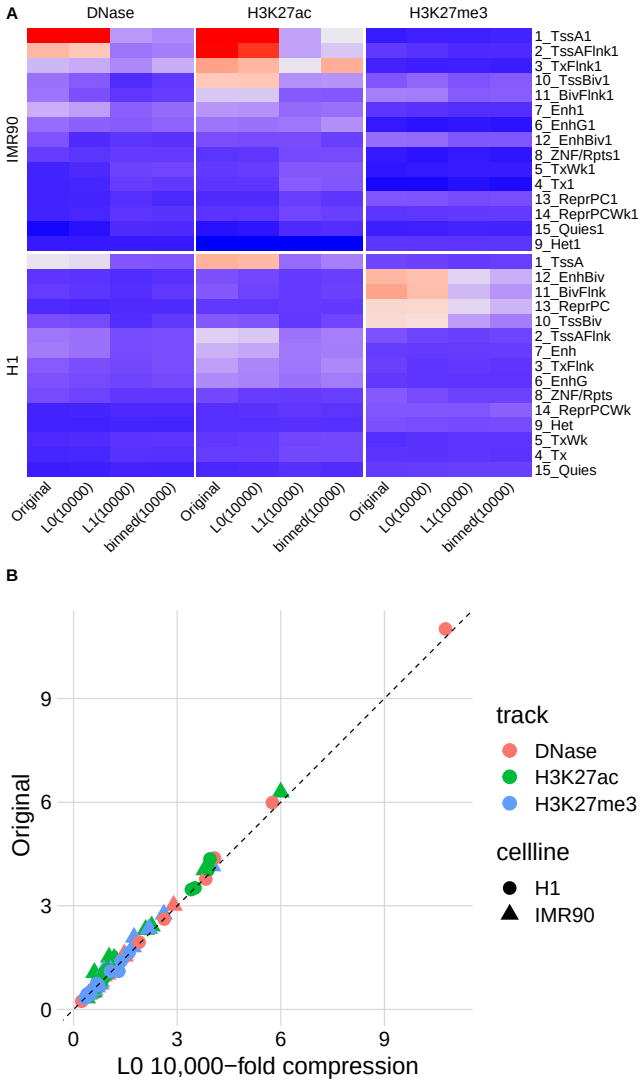

Fig. 4. The ratio of mean signal within ChromHMM clusters to the mean background signal after compressing the signal 10K fold using L_0Poisson (L0), L1 Poisson (L1), and fixed size binning (binned). The results are shown for DNAse, H3K27ac, and H3K27me3 ChIP-seq tracks from H1 and IMR90 cell lines. For all tracks and cell lines, L_0Poisson retains the information content of the signal while L1 Poisson and binning lose either the structure completely or the signal intensity at clusters where the original signal intensifies.

as a count pair of methylated and unmethylated reads. The probability that a base is methylated (the ratio of methylated to total reads) is referred to as β. Importantly, the methylation signal doesn't have natural "peaks" that are the hallmark of other epigenetic tracks such as ATAC-seq and ChIP-seq but is rather dispersed through the genome. Because the read coverage is scattered across the genome, the coverage at any single nucleotide is generally low but variable.

Extending the L_0 segmentation to binomial loss

NGS methylation data can be modeled as a series of draws from a binomial distribution with an unknown success probability. Considering the dual-track input, where y_i^1 is the number of methylated reads and y_i^2 is the number of total reads, it can be derived as follows.

$$\underset{\alpha\in\mathbb{R}^n}{\operatorname{argmax}} \sum_{i=1}^{N} w_i \overbrace{\left(y_i^1\alpha_i - y_i^2 \log\left(1+e^{\alpha_i}\right)\right)}^{\text{Binomial Error}} - \lambda \sum_{i=1}^{N-1} \underbrace{1\left\{\alpha_i - \alpha_{i+1}\right\}}_{L_0\text{ Norm}} \tag{2}$$

Explicitly, modeling the binomial error accurately accounts for read coverage at each location, so that low-coverage segments are more likely to be merged with their neighbors. Secondly, because the L_0 does not depend on the size of the breakpoints, this formulation readily creates single CpG segments if the methylation level is different from that of their neighbors and coverage is sufficient. We demonstrate these two properties of our formulation on a simulated dataset in Fig. 6. In panel A, we simulate a region that has a dramatically different β from its neighborhood, but the value is not reliable because the coverage is low. Segmenting the β values directly using squared error assigns this segment its own region. On the other hand, using the explicit two-track binomial formulation, this region is merged with its neighbors (both segmentations have exactly three segments). In panel B, we simulated a constant methylation region with two single CpG outlier segments, which are readily discovered by L_0 but not by L_1/fused-lasso.

We further illustrate the properties of binomial L_0 segmentation on real WGBS data in H1 and IMR90 cell lines. We consider a region near the POU5F1 locus, which is an important pleuripotency gene and a major differentially methylated region between these two cell lines highlighted in the original WGBS study [28]. In Fig. 7, we consider different segmentations of this region, fixing chromosome-level fold reduction, and we contrast L_0 and L_1. For reference, the differentially methylated region (DMR) of the POU5F1 promoter, discovered with supervised analysis, is shown in purple. Contrasting the L_0 and L_1 formulations, we see that the unsupervised L_0 segmentation is able to identify the promoter DMR in the POU5F1-expressing cell line, H1, at up to 1000-fold compression. The promoter DMR is identified as a feature up to 100-fold reduction in the non-expressing cell line IMR90.

The region around POU5F1 shows a common methylation pattern of expressed genes, where the promoter tends to be hypomethylated while the gene body is hypermethylated. We thus evaluated how well these patterns are preserved globally using compressed representations. Similarly to our previous evaluations of compression for other epigenetic tracks, we evaluate the trade-off between compression level and signal retention, focusing on the chromoHMM states denoted as TssA (TSS) and Tx (transcript). As these are derived from other epigenetic tracks from the same cell line, they mark only the genes that are expressed in each. We find that using raw methylation input, we observe the expected hypomethylation for TssA and hypermethylation for Tx. The pattern is nearly unaltered as we increase the compression level up to 1000-fold for L_0 (see Fig. 8). In contrast, for L_1/fused-lasso, the TssA enrichment begins to decrease at a compression ratio of 100. In summary, we find that the binomial error L_0 segmentation we propose is an effective unsupervised feature reduction methods for NGS methylation data. It naturally accounts for read coverage, can extract very small segments with distinct methylation patterns, and discover salient epigenetic features such as promoters and DMRs in a fully unsupervised manner.

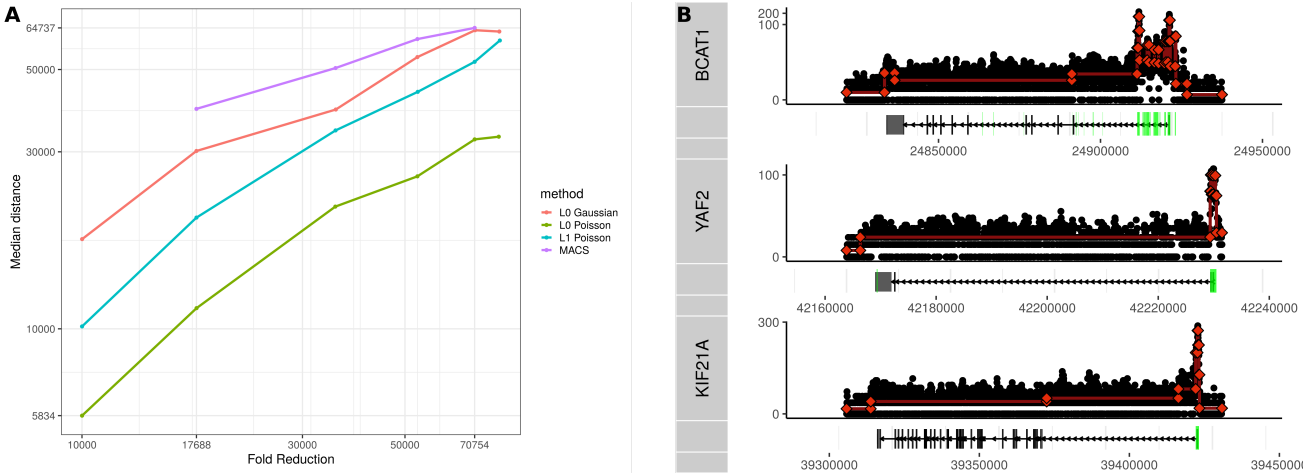

Fig. 5. A: Comparison of the performance of different data reductions approaches on the transcript end site (TES) discovery task. We plot the median absolute distance to TES based on known gene models. The degree of segmentation is adjusted using lambda while for peaks we adjust the FDR cutoff to get a comparable number of segments. **B:** Examples of segmentation tracks for selected transcripts. Green regions indicate MACS peaks. While the Pol II signal is often noisy outside of the promoter region, a subtle drop-off in signal can often be observed towards the end of transcripts. In many cases this change in signal is correctly identified by L0 Poisson segmentation.

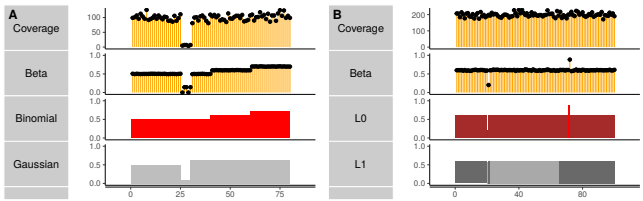

Fig. 6. Binomial L_0segmentation accurately accounts for read coverage and discovers short regions with distinct methylation rates. We compare segmentation formulations on simulated data. Methylation and coverage are shown in yellow. **A:** Squared (gray) and binomial error segmentation (red) with the same number of segments. Binomial error merges low coverage regions with their neighbors. **B:** Comparing L_0 (red) and L_1 (gray). The L_0 penalty is more sensitive to local structure discovering the 2 CpGs simulated with different β.

Conclusion

Epigenetic data refers to a heterogeneous class of assay types that differ in their signal properties. However, all epigenetic data shares a local coherence structure as assay values at nearby nucleotides are determined by a common biological or technical process. This local structure can be leveraged in different ways via algorithms that are specialized for specific data types and data features.

Here, we propose a general computational approach that can exploit local structure in an entirely data-agnostic way. Our approach generalizes a recently proposed highly efficient L_0 segmentation algorithm to probability distributions relevant to modeling epigentic data. We show that the L_0 segmentation efficiently represents epigenetic tracks while retaining many salient features such as peaks, promoters and ChromHMM states, making no assumptions about the underlying signal structure beyond piece-wise constant. This makes our approach suitable for representing diverse track types, including those that have "broad peak" or even complex structure such as Pol II, in a single unifying framework.

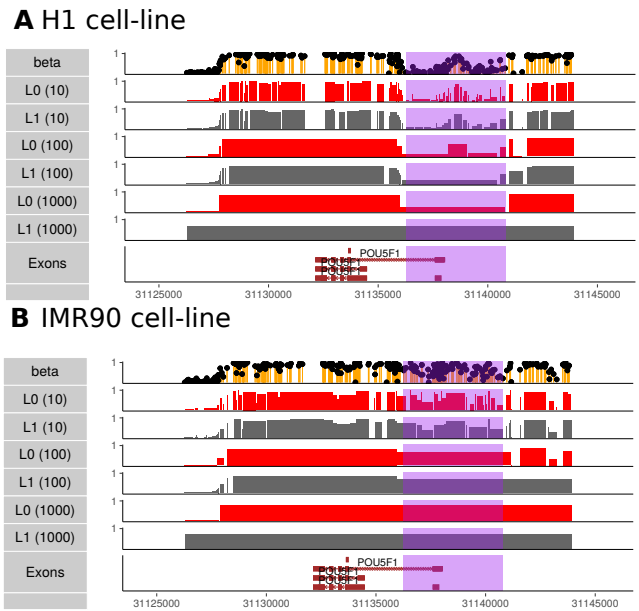

Fig. 7. Comparing L_0 to L_1 with binomial error on real WGBS data with different fold compression computed on the chromosome level. The reported differentially methylated region is highlighted in purple. **A:** POU5F1 expressing cell line, H1. **B:** The non-expressing contrasting cell-line, IMR90. Note that unlike Fig. 6, fold compression is fixed at the chromosome level and consequently the number of segments in this region is not constant across L_0 and L_1.

Methods

Data

All of the datasets used in our analysis are from the IMR90, H1, or K562 cell lines. The example dataset depicted in Fig. 2 is a subset of the K562 BigWigs used in the Segway [29] analysis and was downloaded from `https://noble.gs.washington.edu/proj/segway/2011/bigWig/`. We used the following tracks: "CTCF.Bernstein", "H3K27ac", "H3K27me3", "H3K36me3",

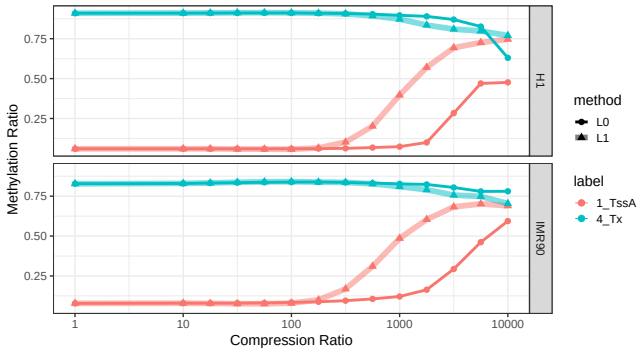

Fig. 8. Comparison of L_0 and L_1 segmentations with binomial error in preserving the characteristics of genomic regions, namely "TssA" (TSS) and "Tx" (Transcript) regions identified by ChromHMM, in two cell lines: H1 and IMR90.

"Pol2_8WG16.Snyder." For our evaluations of ChIP-seq data, we used raw data, MACS peak calls, and ChromHMM results from the H1 and IMR90 cell lines. These datasets are available in the Roadmap data portal at `https://egg2.wustl.edu/roadmap/data/byFileType/alignments/` for raw data, `https://egg2.wustl.edu/roadmap/data/byFileType/peaks/` for MACS peaks, and `https://egg2.wustl.edu/roadmap/data/byFileType/chromhmmSegmentations/` for ChromHMM segmentations. For WGBS, we used the fraction and coverage tracks available at `https://egg2.wustl.edu/roadmap/data/byDataType/dnamethylation/WGBS/FractionalMethylation_bigwig/` and `https://egg2.wustl.edu/roadmap/data/byDataType/dnamethylation/WGBS/ReadCoverage_bigwig/`. For the evaluation of transcription end site (TES) discovery, we used the K562 used Pol II assay BigWig and BAM files provided by the ENCODE data portal with accession code ENCSR000BM.

Software

L_0 segmentation is implemented as an R package with a C++ backend. The embedding of C++ in R is carried out by the Rcpp interface. The implementation is available at `https://github.com/boooooogey/l01segmentation`. Within R, `fusedsegmentation` is the main function that delegates segmentation tasks to various solvers. Through the arguments, users can specify a distribution (Gaussian/Poisson/binomial), a penalty (L_0/L_1), and change the format of the output. Another valuable function is `compressBigwig`, an end-to-end pipeline, that takes a BigWig file as input and generates a BedGraph with segments. Detailed information about the algorithm and our implementation can be found in the appendix.

Binned cross-validation

Cross-validation can be used to find optimal regularization parameters for fused lasso and related segmentation problems [30, 31]. In the standard cross-validation procedure, individual points are left out of the computation and estimated using the value of their segment. The best *lambda* is taken to be the one that achieves minimal test error. If the data are independently sampled from the piece wise constant signal, this procedure will identify the correct λ that recovers the underlying structure without capturing the sampling noise. In the case of epigenetic assays that depend on immunoprecipitation (ChIP), independent sampling conditioned on true signal is not an accurate assumption. Since the protein is pulled down with a DNA fragment that is much larger than its actual occupancy site, regions of the DNA that are not directly bound but are simply near the site are likely to be captured. Moreover, the probability of capture depends continuously on the proximity to the true occupied site, inducing local dependency across adjacent sites and violating independent sampling.

Instead, the value at a single position is best predicted as the average of its immediate neighbors, and using naive cross-validation results in a very large number of segments. To avoid fitting this local structure, we propose a binned cross-validation strategy. Instead of leaving out random or equally spaced elements, the data is divided into windows whose width approximately matches the expected ChIP-seq fragment length (we use 300bp), and the test set is then taken to be equally spaced windows. Furthermore, the test region signal is reconstructed as the average of the left-most point in the adjacent left window and the right-most point in the adjacent right window.

To illustrates, consider a region of length 6 and a cross-validation window of 2. One of the cross-validation test-train splits will be as follows: train: $X_{\text{train}} = X_1, X_2, X_5, X_6$ and test: $X_{\text{test}} = X_3, X_4$. We perform segmentation on X_{train} and we predict $X_3 = X_4 = (\bar{X}_1 + \bar{X}_6)/2$, where $\bar{X}_1, \bar{X}_6$ are the predicted values.

We note that for methylation data, independent sampling is valid and regular cross-validation produced intuitive results.

Transcription end site (TES) analysis

For the transcription end site (TES) discovery task, we utilized the BAM and BigWig files for K562 Pol II signal obtained from ENCODE accession code ENCSR000BMR. These BAM files were merged using the `merge` function from `samtools` and subsequently converted into BigWig files using `bamCoverage`. The L0/1 segmentation framework was applied to the BigWig format, while MACS was employed on the BAM files.

To assess the impact of distinct objective functions and regularizers, we segmented the Pol II signal using the Gaussian and Poisson distributions along with the L0 regularizer, as well as the Poisson distribution with the L1 regularizer. To ensure comparable segmentations across different assumptions, we adjusted the hyperparameter λ. Since broad peaks were not available in the Encode portal we reran the MACS analysis as follows:

```
macs3 callpeak -t rep1_rep2_combined.bam -c
     ctrl1_ctrl2_combined.bam --broad -g hs --broad-cutoff 0.1
      -n broad_peaks_pol2

macs3 callpeak -t rep1_rep2_combined.bam -c
     ctrl1_ctrl2_combined.bam -f BAM -g hs -n test -B -q 0.01
```

Using mostly the default values. The broad and narrow peaks were merged for the evaluation. To enable a fair comparison between MACS peaks and segments, we treated the gaps between peaks as additional segments. During the evaluation, the peaks were sorted based on the adjusted p-values obtained from MACS, and an equivalent number of peaks were selected from the top for comparison.

The test was conducted on a 30 million base pair (bp) segment extracted from the hg38 reference genome, specifically Chromosome 12, starting from the genomic location 8 million bp. The RefSeq gene annotation dataset was utilized as the reference for identifying transcription end sites (TESs). Initially, 807 TESs were identified. Subsequently, additional

filtering was performed based on the Pol II activity at the corresponding transcription start site (TSS) regions. As a result of these steps, 600 TESs remained.

MACS identified 848 peaks within the designated region, which were subsequently divided into four equal parts. In each evaluation step, an additional quarter of the peaks was included. To ensure an equal number of segments for each method, we adjusted the hyperparameter lambda for the L0/1 framework accordingly. The `nearest` function from the GenomicRanges library was employed to calculate the nearest breakpoints for each transcription end site (TES).

Time Comparisons

The evaluation of L0 segmentation using Gaussian and Poisson distributions was performed on the raw read counts derived from BAM files of the same ENCODE experiment as described earlier (accession code ENCSR000BMR). However, for L0 segmentation using the binomial distribution, the evaluation was conducted on the methylation signal obtained from whole-genome bisulfite sequencing (WGBS) data extracted from adult mouse heart tissue, which was part of the dataset published by He et al [32].

Both the methylation and Pol II signals were analyzed using sections of varying lengths. The sections ranged from 100 base pairs (bp) up to 1 million bp, with each step increasing the length by a factor of 10. These sections were extracted from Chromosome 10, starting from the 10 million bp position.

Throughout the experiment, a hyperparameter lambda of 50 was arbitrarily set. It's worth noting that the value of this hyperparameter does not impact the runtime of the framework.

To compare the CROCS, PELT, and L0 frameworks, we utilized the same ENCODE experiment focusing on POL II. Our tests were conducted on Chromosome 1, using sections of varying lengths. The sections started at 250 base pairs (bp) and ended at 250,000 bp, with each step increasing the length by a factor of 10.

For CROCS, we downloaded the implementation from the GitHub repository `aLiehrmann/CROCS` and used the default parameter values during the experiments.

Regarding PELT, we did not have access to the exact implementation used by Yokoyama et al [13]. However, we were able to find an implementation in the `changepoint` library. We set the penalty type to `"Manual"`. Increasing the penalty hyperparameter leads to a more extensive pruning process, which results in longer execution times. Considering our focus on biologically meaningful regions, we determined that a penalty parameter value of 50000 produced meaningful peaks for POL II in a small example. Consequently we decided to set the penalty parameter of PELT to this value throughout the time comparison study.

To ensure comparability between the L0 segmentation results and PELT, we set the hyperparameter lambda for L0 in such a way that it resulted in the same number of breakpoints as obtained by PELT.

Competing interests

No competing interest is declared.

Author contributions statement

Must include all authors, identified by initials, for example: M.C. and T.B. conceived the approach, developed algorithms, analyzed the results, and wrote the paper.

References

1. Ian Dunham, Anshul Kundaje, and et al. An integrated encyclopedia of DNA elements in the human genome. *Nature*, 489(7414):57–74, September 2012. Number: 7414 Publisher: Nature Publishing Group.
2. Anshul Kundaje and et al. Integrative analysis of 111 reference human epigenomes. *Nature*, 518:317–330, February 2015.
3. Yong Zhang, Tao Liu, Clifford A. Meyer, Jérôme Eeckhoute, David S. Johnson, Bradley E. Bernstein, Chad Nusbaum, Richard M. Myers, Myles Brown, Wei Li, and X. Shirley Liu. Model-based Analysis of ChIP-Seq (MACS). *Genome Biol.*, 9(9):1–9, November 2008.
4. macs3 project. MACS, December 2022. [Online; accessed 15. Dec. 2022].
5. Kasper D. Hansen, Benjamin Langmead, and Rafael A. Irizarry. BSmooth: from whole genome bisulfite sequencing reads to differentially methylated regions. *Genome Biology*, 13(10):R83, October 2012.
6. Yutaka Saito, Junko Tsuji, and Toutai Mituyama. Bisulfighter: accurate detection of methylated cytosines and differentially methylated regions. *Nucleic Acids Research*, 42(6):e45, April 2014.
7. Xiaoqing Yu and Shuying Sun. HMM-DM: identifying differentially methylated regions using a hidden Markov model. *Statistical Applications in Genetics and Molecular Biology*, 15(1):69–81, March 2016.
8. Altuna Akalin, Matthias Kormaksson, Sheng Li, Francine E. Garrett-Bakelman, Maria E. Figueroa, Ari Melnick, and Christopher E. Mason. methylKit: a comprehensive R package for the analysis of genome-wide DNA methylation profiles. *Genome Biology*, 13(10):R87, October 2012.
9. Yongseok Park, Maria E. Figueroa, Laura S. Rozek, and Maureen A. Sartor. MethylSig: a whole genome DNA methylation analysis pipeline. *Bioinformatics*, 30(17):2414–2422, September 2014.
10. Arnaud Liehrmann, Guillem Rigaill, and Toby Dylan Hocking. Increased peak detection accuracy in over-dispersed chip-seq data with supervised segmentation models. *BMC bioinformatics*, 22(1):1–18, 2021.
11. Vincent Runge, Toby Dylan Hocking, Gaetano Romano, Fatemeh Afghah, Paul Fearnhead, and Guillem Rigaill. gfpop: an r package for univariate graph-constrained change-point detection. *arXiv preprint arXiv:2002.03646*, 2020.
12. Rebecca Killick, Paul Fearnhead, and Idris A Eckley. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598, 2012.
13. Takao Yokoyama, Fumihito Miura, Hiromitsu Araki, Kohji Okamura, and Takashi Ito. Changepoint detection in base-resolution methylome data reveals a robust signature of methylated domain landscape. *BMC genomics*, 16(1):1–10, 2015.
14. Sabrina Herrmann, Holger Schwender, Katja Ickstadt, and Peter Müller. A bayesian changepoint analysis of chip-seq

data of lamin b. *Biochimica et Biophysica Acta (BBA)-Proteins and Proteomics*, 1844(1):138–144, 2014.

15. Mengjie Chen, Haifan Lin, and Hongyu Zhao. Change point analysis of histone modifications reveals epigenetic blocks linking to physical domains. *The annals of applied statistics*, 10(1):506, 2016.
16. Haipeng Xing, Yifan Mo, Will Liao, and Michael Q Zhang. Genome-wide localization of protein-dna binding and histone modification by a bayesian change-point method with chip-seq data. *PLoS computational biology*, 8(7):e1002613, 2012.
17. Nicholas A. Johnson. A Dynamic Programming Algorithm for the Fused Lasso and L0-Segmentation. *Journal of Computational and Graphical Statistics*, 22(2):246–260, 2013.
18. Shuying Sun and Xiaoqing Yu. HMM-Fisher: identifying differential methylation using a hidden Markov model and Fisher's exact test. *Statistical Applications in Genetics and Molecular Biology*, 15(1):55–67, March 2016. Publisher: De Gruyter.
19. Boying Gong and Elizabeth Purdom. Methcp: differentially methylated region detection with change point models. *Journal of Computational Biology*, 27(4):458–471, 2020.
20. Jason Ernst and Manolis Kellis. Chromatin state discovery and genome annotation with ChromHMM. *Nat. Protoc.*, 12(12):2478, December 2017.
21. Guoqiang Li, Yaping Liu, Yanxiao Zhang, Naoki Kubo, Miao Yu, Rongxin Fang, Manolis Kellis, and Bing Ren. Joint profiling of DNA methylation and chromatin architecture in single cells. *Nat. Methods*, 16(10):991, October 2019.
22. Sarah Dedeurwaerder, Matthieu Defrance, Martin Bizet, Emilie Calonne, Gianluca Bontempi, and François Fuks. A comprehensive overview of Infinium HumanMethylation450 data processing. *Briefings Bioinf.*, 15(6):929, November 2014.
23. Zelin Jin and Yun Liu. DNA methylation in human diseases. *Genes & Diseases*, 5(1):1, March 2018.
24. Tessa Bergsma and Ekaterina Rogaeva. DNA Methylation Clocks and Their Predictive Capacity for Aging. *Neuroscience Insights*, 15, 2020.
25. Morgan Levine, Ross A. McDevitt, Margarita Meer, Kathy Perdue, Andrea Di Francesco, Theresa Meade, Colin Farrell, Kyra Thrush, Meng Wang, Christopher Dunn, Matteo Pellegrini, Rafael de Cabo, and Luigi Ferrucci. A rat epigenetic clock recapitulates phenotypic aging and co-localizes with heterochromatin. *eLife*, November 2020.
26. Todd R. Robeck, Zhe Fei, Ake T. Lu, Amin Haghani, Eve Jourdain, Joseph A. Zoller, Caesar Z. Li, Karen J. Steinman, Stacy DiRocco, Todd Schmitt, Steve Osborn, Bill Van Bonn, Etsuko Katsumata, June Mergl, Javier Almunia, Magdalena Rodriguez, Martin Haulena, Christopher Dold, and Steve Horvath. Multi-species and multi-tissue methylation clocks for age estimation in toothed whales and dolphins. *Commun. Biol.*, 4(642):1–11, May 2021.
27. Mike Thompson, Brian L. Hill, Nadav Rakocz, Jeffrey N. Chiang, , Sriram Sankararaman, Ira Hofer, Maxime Cannesson, Noah Zaitlen, and Eran Halperin. Methylation risk scores are associated with a collection of phenotypes within electronic health record systems. *medRxiv*, page 2022.02.07.22270047, February 2022.
28. Ryan Lister, Mattia Pelizzola, Robert H. Dowen, R. David Hawkins, Gary Hon, Julian Tonti-Filippini, Joseph R. Nery, Leonard Lee, Zhen Ye, Que-Minh Ngo, Lee Edsall, Jessica Antosiewicz-Bourget, Ron Stewart, Victor Ruotti, A. Harvey Millar, James A. Thomson, Bing Ren, and Joseph R. Ecker. Human DNA methylomes at base resolution show widespread epigenomic differences. *Nature*, 462:315–322, November 2009.
29. Michael M. Hoffman, Orion J. Buske, Jie Wang, Zhiping Weng, Jeff A. Bilmes, and William Stafford Noble. Unsupervised pattern discovery in human chromatin structure through genomic segmentation. *Nat. Methods*, 9:473–476, May 2012.
30. Daniela M. Witten, Robert Tibshirani, and Trevor Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics (Oxford, England)*, 10(3):515, July 2009.
31. Sean Jewell and Daniela Witten. EXACT SPIKE TRAIN INFERENCE VIA 0 OPTIMIZATION. *annals of applied statistics*, 12(4):2457, December 2018.
32. Yupeng He, Manoj Hariharan, David U Gorkin, Diane E Dickel, Chongyuan Luo, Rosa G Castanon, Joseph R Nery, Ah Young Lee, Yuan Zhao, Hui Huang, et al. Spatiotemporal dna methylome dynamics of the developing mouse fetus. *Nature*, 583(7818):752–759, 2020.