Inferring gene-regulatory networks using epigenomic priors

Thomas E. Bartlett[1*], Melodie Li[1], Chenyu Song[1], Yuche Gao[1], and Qiulin Huang[2,3]

[1]Department of Statistical Science, University College London, London, U.K.
[2]Human Embryo and Stem Cell Laboratory, The Francis Crick Institute, London, U.K.
[3]Department of Physiology, Development and Neuroscience, The Centre for Trophoblast Research, University of Cambridge, Cambridge, U.K.
[*]thomas.bartlett.10@ucl.ac.uk

Abstract

We show improved accuracy of inference of gene-regulatory network (GRN) structure, resulting from the use of an epigenomic prior network. We do so by re-analysing 12 datasets from 12 different studies, including cells from human embryos, healthy adult tissue, and cancer, including single-cell mRNA sequencing data, DNA methylation (DNAme) data, chromatin accessibility data, and histone modification data. We find that DNAme data are very effective for inferring the epigenomic prior network, recapitulating known epigenomic network structure found previously from chromatin accessibility data. Furthermore, we find that inferring the epigenomic prior network from DNAme data reveals candidate TF cis-regulations for at least eight times as many genes, when compared with chromatin accessibility data. When our proposed methodology is applied to real datasets from human embryonic development and from women at risk of breast cancer, we find patterns of differential cis-regulation that are in line with expectations under appropriate biological models, and that can be used to identify pre-cancerous epigenomic changes.

Keywords

DNA methylation; gene-regulatory networks; human embryonic development; breast cancer biomarkers.

1 Introduction

Gene-regulatory network (GRN) models have been shown to be effective at finding specific panels of transcription factors (TFs) that are associated with cell-type identities and with progression of certain diseases (see for example [1–3]). These successful methods of GRN inference are typically based on the target-gene approach [4–6] (Equation 1), which can find important structure in 'omic networks such as feedback loops [7]. Here, we propose a novel improvement to target-gene GRN inference methodology that builds on our previous work to infer networks from epigenomic data such as DNA methylation (DNAme) data [8, 9], and using advanced regression methods with scRNA-seq (single-cell RNA-seq) and multimodal data [10–12]. Our proposed methodology infers the GRN using advanced regression methods with scRNA-seq (gene-expression) data *following* estimation of an epigenomic prior network. The sequence of these operations is crucial. We show that the accuracy of GRN inference from gene expression data is much improved when using an epigenomic prior network that has been inferred in advance. The epigenomic prior network provides a list of candidate TF regulators that is specific to each target-gene, for subsequent refinement with expression data. This contrasts with alternative strategies, in which the GRN is estimated by combining networks that have been inferred separately and independently from gene expression data and from chromatin accessibility data [1, 2, 5]. Importantly, with our proposed methodology the gene-expression data can originate from a completely independent source to the data that is used to infer the epigenomic prior network. This enables novel analyses specific to tissue and cell-type, by making use of the wealth of genomic and epigenomic data that is now publicly available [13].

A gene-regulatory network (GRN) represents all the gene targets of every TF in a specific cell type. Hence, GRNs are important summaries of the genome-wide gene regulation profile that characterises a particular cellular identity. GRNs have been inferred successfully using methods based on random forests (such as GENIE3 / GRNboost2 [5, 6] as part of the SCENIC / SCENIC+ pipelines [1–3]), as well as alternative approaches based on mutual information [12, 14–17] and recent extensions [18]. However, a genomic network that has been inferred only from gene expression data (such as scRNA-seq data) is more accurately referred to as a gene co-expression network, unless data are included on the interactions between gene-products and DNA [10, 19, 20]. Hence, in order to infer GRNs, data must be included in the inference procedure that allows more precise inference about the interaction of the gene product of the regulating gene (i.e., a TF) with the DNA near the regulated gene (i.e., *cis-regulation*) [1, 3, 11, 12]. Such data often combine the DNA location of cis-regulatory TF binding sites with

epigenomic measurements at those locations such as from chromatin accessibility as assessed by ATAC-seq [21]. However, alternative epigenomic data-sources include DNAme data, and ChIP-seq data for histone modifications such as H3K27ac, H3K4me3, and K3K27me3.

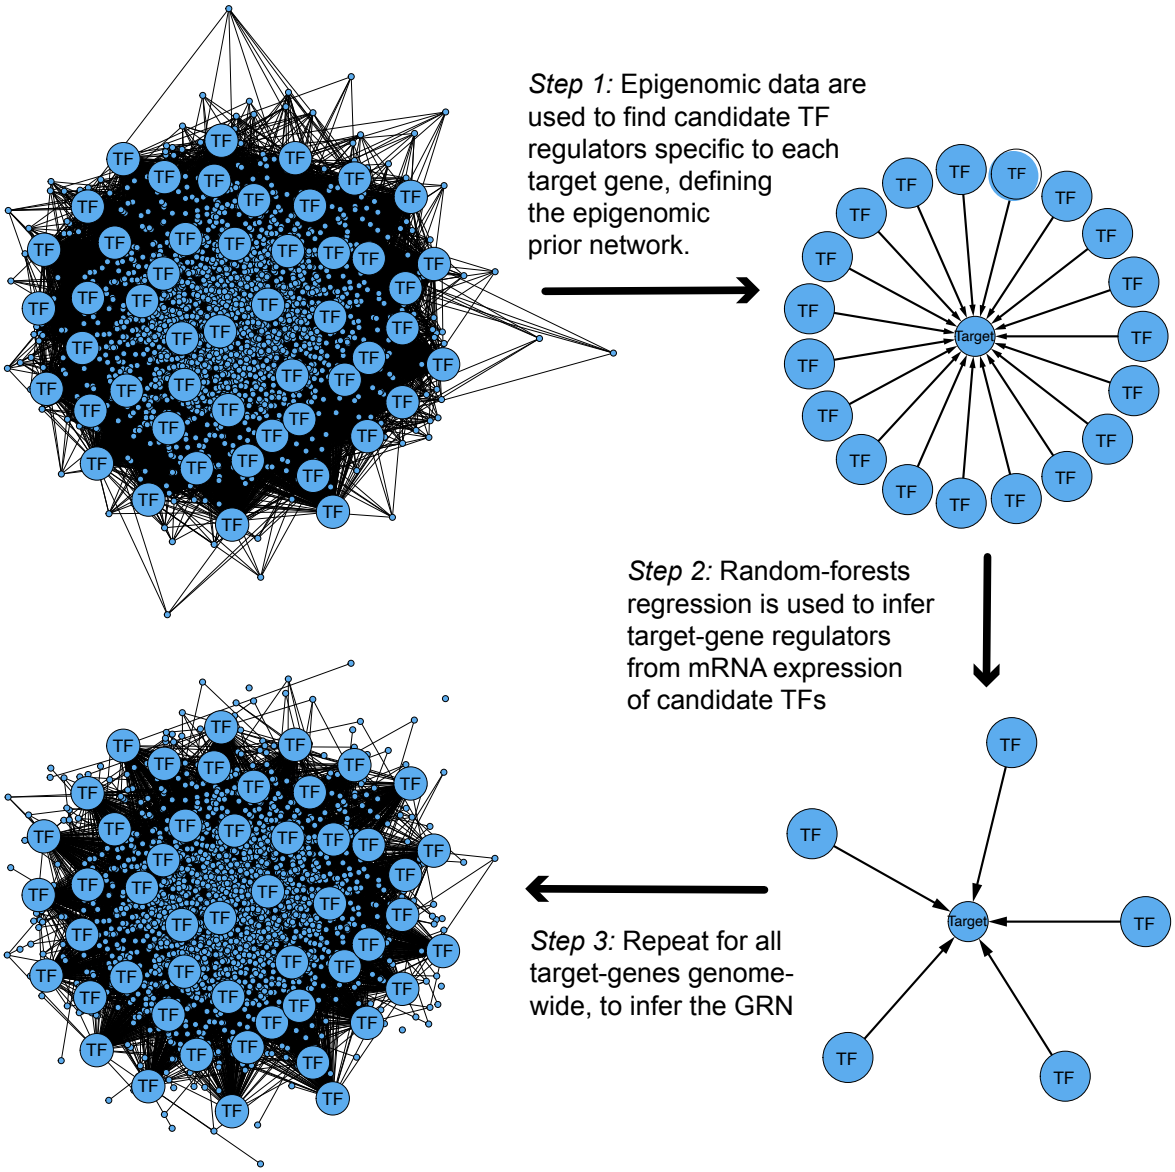

Figure 1: *Proposed gene-regulatory network (GRN) inference procedure, for both real data and simulation study. Step 1: Epigenomic data (such as ATAC-seq / LiCAT-seq, H3K27ac, DNAme, data) are used to infer the prior network structure, by analysing epigenomic patterns at cis-regulatory TF binding sites [22, 23]. This provides a set of candidate TF regulators specific to each target-gene, based on whether a candidate gene-regulation is physically realistic. Step 2: An advanced regression method such as random-forests is used to infer the cis-regulators of the target-gene from the specific set of candidate TFs defined in the epigenomic prior network. Step 3: The global gene cis-regulatory network is constructed, by combining the local networks that were inferred around each target-gene [4] in the previous steps.*

In the target-gene approach to GRN inference [4–6], an advanced regression method is used to model the response x_i (the expression-level of gene i, Figure 1), as:

$$x_i = f_i(\mathbf{x}_{\mathbf{PA}^{\text{prior}}(i)}, e_i), \tag{1}$$

where $\mathbf{PA}^{\text{prior}}(i)$ are the candidate regulators of gene i, $\mathbf{x}_{\mathbf{PA}^{\text{prior}}(i)}$ is the vector of expression levels of these candidate regulators, f_i is a (possibly non-linear) function, and e_i is an error or noise term. In our proposed methodology, the epigenomic prior network consists of the candidate regulators of all p genes genome-wide, $\mathbf{PA}^{\text{prior}}(i)$, $i \in \{1, ..., p\}$, that are inferred from epigenomic data in combination with TF motif-DNA sequence matching (Figure 1). In our proposed methodology, we first estimate the set of candidate TF regulators $\mathbf{PA}^{\text{prior}}(i)$ separately for each target-gene i from epigenomic data. This is part of an inference procedure where the second step uses advanced regression with gene-expression data to infer the regulators of each i from the candidates provided by $\mathbf{PA}^{\text{prior}}(i)$ (Equation 1, Figure 1). State-of-the-art methods for target-gene GRN inference such as GENIE3 / GRNboost2 [5,6] (as part of the SCENIC / SCENIC+ pipelines [1–3]) do not use candidate sets of regulator TFs $\mathbf{PA}^{\text{prior}}(i)$ that are specific to target-gene i. This makes it more difficult for the advanced regression method to estimate f_i (Equation 1), because the non-target-specific set of candidate TFs must be much larger (of the order of 100s or 1000s) than the target-specific sets (of the order of 10s of TFs). Hence, we use the same type of advanced regression (random-forests) as algorithms such as GENIE3 and GRNboost2 do, but with the novel improvement that in the advanced regression step to estimate f_i, we use much smaller sets of candidate TFs $\mathbf{PA}^{\text{prior}}(i)$ that are specific to each target-gene i. These target-specific candidate TF regulators $\mathbf{PA}^{\text{prior}}(i)$, $i \in \{1, ..., p\}$, are encoded in the epigenomic prior network.

It is well known that changes in DNAme patterns precede cancer formation [24, 25]: these pre-cancerous changes are sometimes called 'epigenomic field-defects' [24, 26] or 'epigenetic field cancerisation' [27]. We have previously found that that unique epigenomic patterns distinguish cancer [28], and can be used to predict clinical outcomes [29], and it is well established that DNAme patterns enable different tissues and cell types to be distinguished from one another [30]. Changes in DNAme rates at regulatory DNA regions enable changes in gene regulation [31] and therefore allow changes of phenotype [32], motivating the methodology we propose here to systematically analyse genome-wide DNAme rates at putative binding sites of specific TFs. Recently, it has become clear that a large amount of the observed variation in DNAme patterns in bulk-tissue genomic data can be explained by the variation in the proportion of cells from different lineages that comprise the tissue sample [33]. Hence bulk-tissue DNAme data contain important information about the cell type composition of the tissue, and this finding can also be used as a clinical biomarker, by tracking cells-of-origin of cancer [27, 34]. In the methodology proposed here, we infer GRNs using DNAme data from specific cell types (rather than from a mixture of cells of different lineages), to allow inferences about the functional effects of DNAme changes at cis-regulatory TF binding sites. We use these cell-type specific DNAme data to infer genomic networks, showing how DNAme data can be used to infer the structure of the epigenomic prior network.

It has previously been established that DNAme data are very effective for GRN inference in combination with gene expression data [35]. We have found that an epigenomic prior network (for example based on DNAme data) is essential to achieve the most accurate inference from the advanced regression methods applied to gene expression data such as scRNA-seq data (Figures 1, 2, 3, and 7, and Table 1). We show that the epigenomic prior networks we infer recapitulate well-validated TF cis-regulations of target-genes obtained using the widely used and well validated software packages pycistarget [1, 3, 23] and RGT-HINT [36]. We then infer GRNs from gene-expression (scRNA-seq) data in combination with these epigenomic prior networks, in the contexts of data from human embryonic development, and from studies of women at risk of breast cancer. We find GRNs characteristic of different lineage choices and disease at-risk states, with patterns of differential gene-expression that are in line with expectations under appropriate biological models. However, we note that novel findings from computational GRN inference remain hypotheses until they are experimentally validated. Therefore to give confidence that our proposed methodology leads to verifiable inferences, we use the most significant inferences from our methodology as case studies, showing that they correspond to existing knowledge that is well supported by previous studies.

2 Results

2.1 Using prior network structure from epigenomic data improves performance of regulatory network inference from gene expression data

We found that the epigenomic prior network allows the advanced regression method (Equation 1) to make more accurate predictions of gene regulations, in the context of biologically-realistic data generated from known network structure (Figure 2). When there are fewer candidate regulators of a target-gene, we found improved accuracy (assessed by AUC) identifying known gene regulations, as shown in Figure 3. These results are based on data generated from known network structure, that was simulated using realistic differential-equations models of chemical reaction kinetics [37] (Section 4). We hence use the AUC statistic to test how well the network edge strengths found by advanced regression, can predict the known ground-truth network structure (that is not otherwise used in the network inference). It can be seen from Figure 3 that there is an increase in accuracy of GRN inference

resulting from the reduction in the number of candidate TF regulators that the advanced regression method must choose from, from the order of 100s or 1000s down to the order of 10s. This reduction in number of candidate regulators is enabled by the epigenomic prior network.

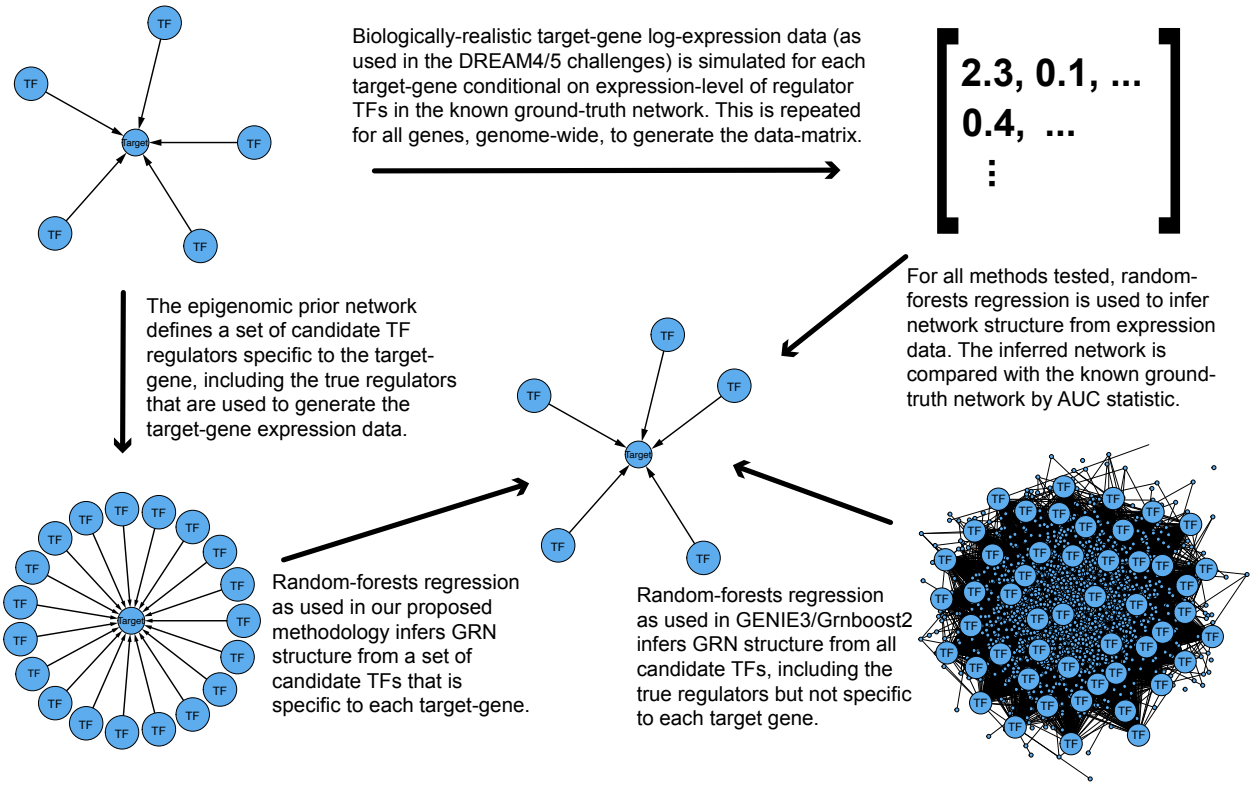

Figure 2: *Biologically-realistic gene-expression data (as used in the DREAM4/5 challenge) was generated based on known ground-truth network structure. The prior network provides target-specific lists of candidate TF regulators, which were generated by combining the list of true regulators defined by the ground-truth network, with a random sample of other regulators found in the ground-truth network. Random-forests regression was then applied in two ways. Firstly, to the target-specific lists of candidate regulators defined in the prior network, as in our proposed methodology (Figure 1). Secondly, to the full list of all regulators in the prior network, i.e. not target-specific as required in GENIE3/Grnboost2. The findings are shown in Figure 3*

For each target-gene, the epigenomic prior network defines a list of candidate TF regulators of the target-gene that are specific to that target-gene. This is based on whether the regulation of target by TF is physically realistic. The target-gene specific list of candidate regulators is typically a few 10s of TFs, making it much easier for the advanced regression method to identify the correct TF regulators of the target-gene (Figure 3a), than it would be from a much longer list of candidate TFs that is not target-specific, as required by GENIE3 / GRNboost2 and similar [1–3, 5, 6]. With existing methods such as GENIE3 / GRNboost2, the full list of TFs (typically of the order of 100s or 1000s of TFs) must be specified as candidate regulators of each target-gene, meaning that it is much more challenging for those methods to find the true TF regulators of the target-gene (Figure 3b). Hence, the epigenomic prior network allows greater accuracy of GRN inference from advanced regression methods such as random forests [38], when used with single-cell gene expression data in a target-gene GRN inference procedure [4].

2.2 Recapitulating known epigenomic network structure using DNA methylation data

For the human embryonic development LiCAT-seq (low-input chromatin accessibility) dataset [39], we first verified whether TF-motif cis-regulations that were found using RGT-HINT [36] in previous work [12] would also be discovered from the same LiCAT-seq data by our proposed modified-lever-cistarget (MLC) method for epigenomic prior network inference (Section 4). We found 170 TFs that were present in both analyses, and found that the corresponding regulatory regions overlapped significantly for 168 out of these 170 TFs (Figure 4a). We also applied pycistarget [1, 3, 23] to the same LiCAT-seq human embryonic development dataset [39], finding 124 signficant TF motifs: a large number of these 124 motifs found by pycistarget were also found to be significant according to the

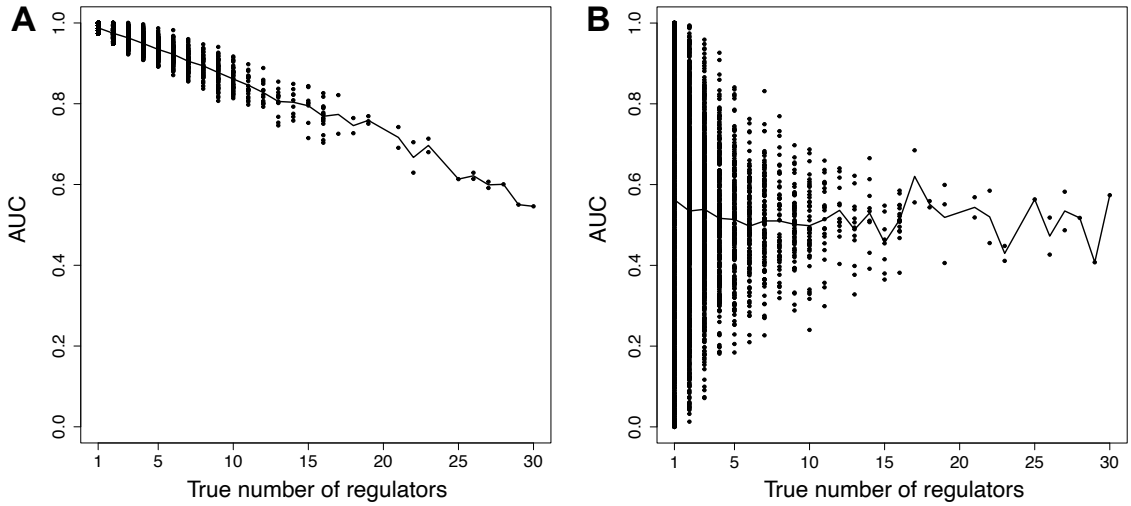

Figure 3: *Results from GRN inference using random-forests regression with biologically-realistic gene-expression data generated from known ground-truth network structure. (a) Random-forests target-gene GRN inference implemented as in our proposed methodology (Figure 1). The prior network structure provides a target-specific set of candidate regulators, from which random-forests regression chooses. (b) Random-forests target-gene GRN inference, implemented with GENIE3. Random forests regression infers the target regulators by choosing from all candidate regulators that are found in the prior network, i.e., not from a list that is specific to each target gene.*

MLC method (Figure 4b). Hence, we were able to validate the TF-motif cis-regulatory inferences found by MLC from the LiCAT-seq human embryonic development dataset [39], in comparison with TF-motif cis-regulations that were found from same dataset using RGT-HINT [36] in previous work [12], and also in comparison with inferences from the pycistarget method which is part of the SCENIC+ pipeline [3].

Applying the MLC method to the LiCAT-seq human embryonic development dataset [39] allows an epigenomic prior network to be inferred. Equivalently, applying our proposed MethCorTarget methodology (Section 4) to the DNAme human embryonic development dataset [40] allows an epigenomic prior network to be inferred. These epigenomic prior networks inferred from different epigenomic data sources are expected to contain overlapping and similar information. We verified this by testing the overlap between the TF motif-gene cis-regulations inferred by MLC as applied to the LiCAT-seq human embryonic development dataset [39], and the cis-regulations inferred by MethCorTarget applied to the DNAme human embryonic development dataset [40]: we found a highly significant overlap, odds-ratio OR $= 13.3$ $(12.9 - 13.8)$ ($p < 0.001$, Fisher's exact test). This large overlap between the lists of cis-regulations inferred from the different data-sources illustrates the interchangeability between the epigenomic prior networks obtained for the same cell type from DNAme data and from chromatin accessibility data. We also note that epigenomic prior network structure inferred from DNAme may be more useful in certain applications, because it provides information about the regulation of more than eight times as many genes (23575 regulated genes found by MethCorTarget from DNAme data, compared to 2650 for MLC from LiCAT-seq data).

2.3 Inferring GRNs in epiblast and primitive endoderm cells of the human embryo

Epiblast (Epi) and primitive endoderm (PE, or hypoblast) cells are the main cell types of the inner cell mass (ICM) of the blastocyst of the human embryo at 5-7dpf (days post fertilisation) [41]. We have previously used ICM chromatin accessibility data to inform inference of GRNs from scRNA-seq data in Epi and PE [12]. We found that our epigenomic prior network obtained from DNAme data overlaps strongly with that obtained from chromatin accessibility data, but covers a much larger number of target-genes. We therefore proceeded to use the epigenomic prior network obtained from DNAme data as a valid epigenomic prior network to improve GRN inference from the human embryonic development scRNA-seq dataset [42]. We inferred GRNs from scRNA-seq data as described (Figure 1 and Section 4) for both Epi and PE cell types, using the scRNA-seq dataset cell type definitions shown in Figure S1 and Figure S2, and using the epigenomic prior network obtained from the DNAme human embryonic development dataset (ICM cells) [40].

An example of cis-regulations inferred as part of the GRN for Epi and PE is shown in Figure 5. The GRN was inferred by filtering candidate cis-regulations found as part of the epigenomic prior network from DNAme data for ICM cells, by variable selection using advanced regression methods and scRNA-seq data for Epi and PE (according to the procedure outlined in Figure 1). Figure 5 shows DNAme rate profiles for the 103 ICM cells (which include Epi cells and PE cells) in the DNAme human embryonic development dataset [40] for the *SOX17* gene. The *SOX17* gene is a well known marker gene of PE cells in contrast with Epi cells [44], and hence we have used it as an ex-

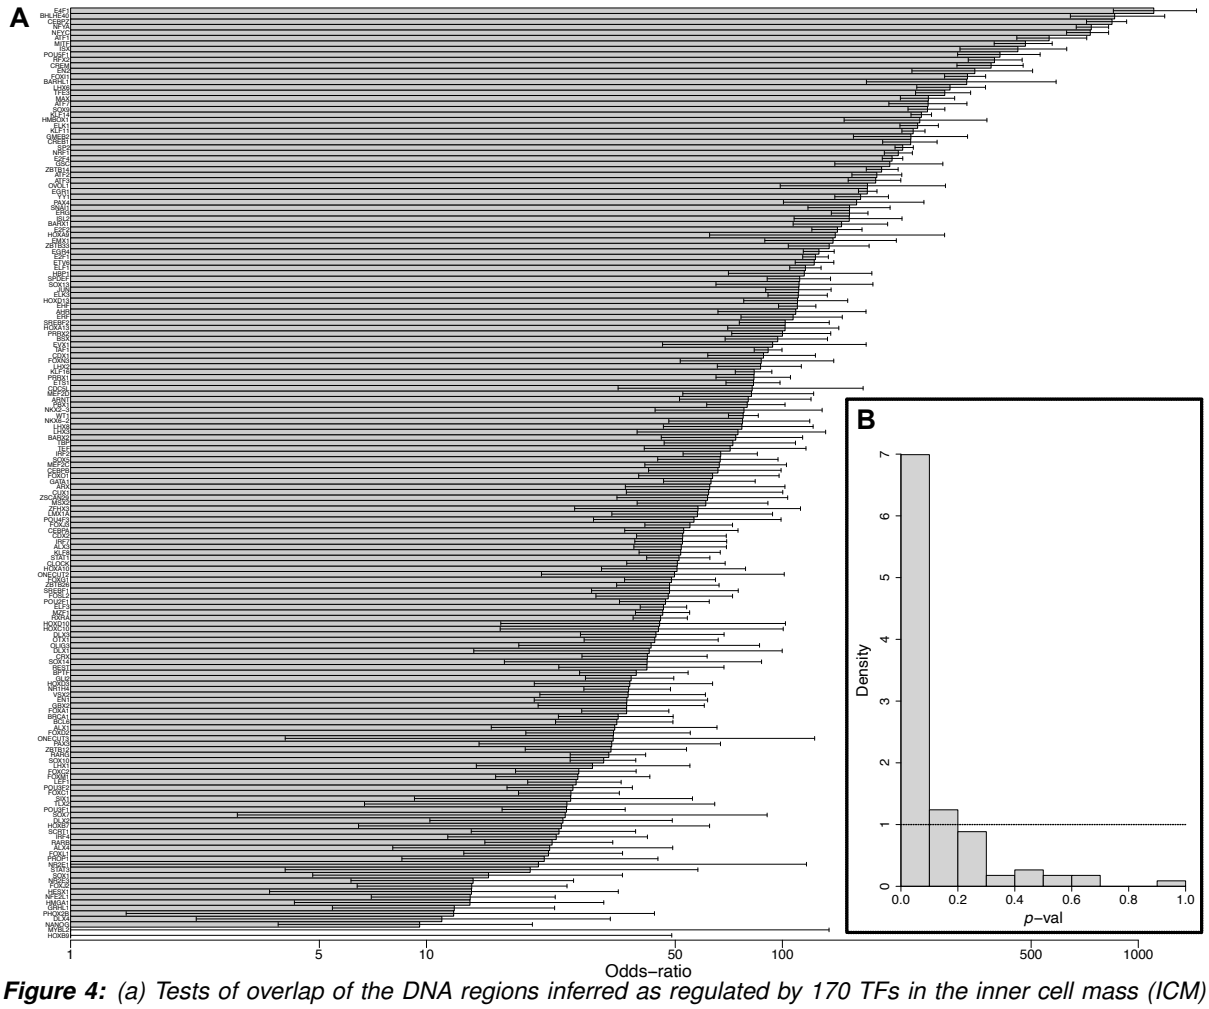

Figure 4: *(a) Tests of overlap of the DNA regions inferred as regulated by 170 TFs in the inner cell mass (ICM) according to both the modified-lever-cistarget (MLC) method and the RGT-HINT method [36] (as published previously [12]). Bars show odds-ratios (on a logarithmic scale), with whiskers showing 95% C.I.s. The corresponding significance p-vals (Fisher's exact test) are provided in Table S1. (b) The 124 motifs inferred when the pycistarget software is applied to the same LiCAT-seq data are are tested for significance according to the MLC method (z-test). The horizontal dashed line at density of 1 indicates the expected density for a uniform distribution of p-values, which would be expected if there were no correspondence between the TF-motif inferences from pycistarget and MLC for this LiCAT-seq dataset. LiCAT-seq is low-input chromatin accessibility data.*

ample here. In Figure 5, changes in the DNAme profile are visible in regions corresponding to previously-published and validated LiCAT-seq peaks [12, 39], as well as around regulatory regions corresponding to the motifs of candidate TF-regulators of *SOX17*. Figure 5b also shows how in epiblast cells, the epiblast marker gene *NANOG* (Figure S1 and Figure S3) negatively regulates *SOX17* (Figure S2 and Figure S3), which is consistent with *NANOG* maintaining Epi identity by repressing PE identity [43]. Also interestingly, Figure 5 illustrates how the DNAme rates at these regulatory regions are highly variable, a feature we note is also of interest in understanding cancer formation and cancer predisposition [28, 29]. There are well established similarities between gene-regulatory processes that have been observed in embryonic development and in cancer [45], and we hypothesise that phenotypic plasticity is represented by the variability in methylation rates that is visible in Figure 5 at CpG loci around the motifs of the candidate TF cis-regulators of *SOX17*.

2.4 Differential regulation in epiblast and primitive endoderm cells

Differentially regulating TFs were identified as follows. For each TF, the differences in regulatory strength (measured by random-forests importance gain) were tested in Epi compared to PE, for all genes the TF can regulate according to the epigenomic prior network. Figure 6 shows that the differential regulation in Epi and PE is characterised by consistent upregulation of genes in PE, whereas in Epi there is significantly more downregulation of genes ($p < 0.001$, two-sample z-test of proportions). This follows from the differential expression analysis (comparing Epi with PE, Section 4), where the majority of genes were found to be upregulated in PE, relative to

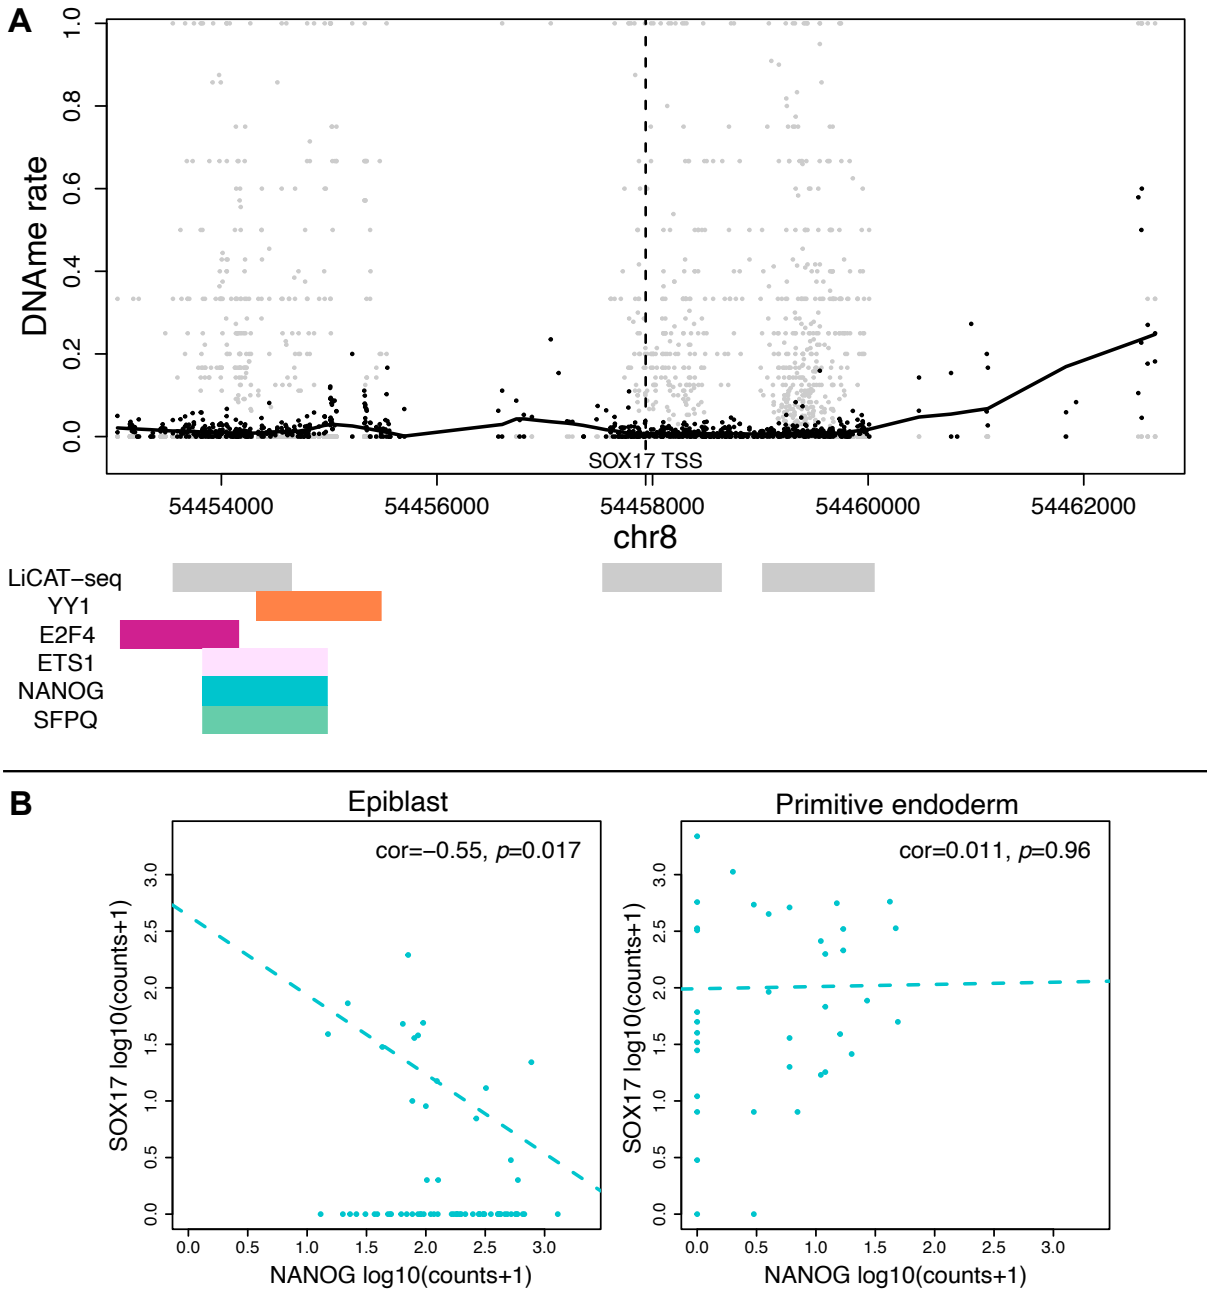

Figure 5: *(a) The DNAme profile of the cis-regulatory region (+/- 5kb) around the TSS (transcriptional start site) of the SOX17 gene (used as a marker for primitive endoderm cells in Figure S1, Figure S2, and Figure S3). Also shown are regions corresponding to previously-published and validated LiCAT-seq peaks [12, 39], and regulatory regions corresponding to motifs of candidate TF regulators inferred as part of the GRN. Mean DNAme rates are shown in black, and those for individual cells are shown in grey; non-linear trend-line is calculated with LOESS regression. (b) Scatter plots of log expression levels of SOX17 and one of its regulators NANOG [43] in Epi and PE.*

Epi. We note that this is a feature of differential gene regulation that we would expect to find if the epiblast can be thought of as the default cell type at the lineage bifurcation of the epiblast and hypoblast (primitive endoderm), as has been shown recently in human embyros [46]. In this case, specification of the hypoblast / primitive endoderm should be characterised by the activation in PE of a characteristic GRN relative to Epi, such as the differential regulation network shown in Figure 6. We also tested enrichment of the TF targets of the differentially regulating transcription factors shown in Figure 6, by comparison with 713 epiblast gene-sets defined and validated previously

by domain-experts [12]. We confirmed that there was significantly more enrichment of these epiblast gene-sets amongst TF targets differentially regulated in epiblast compared to primitive endoderm ($p = 0.05$, *t*-test), i.e., when comparing the enrichment of TF targets in Figure 6a with the enrichment of those in Figure 6b.

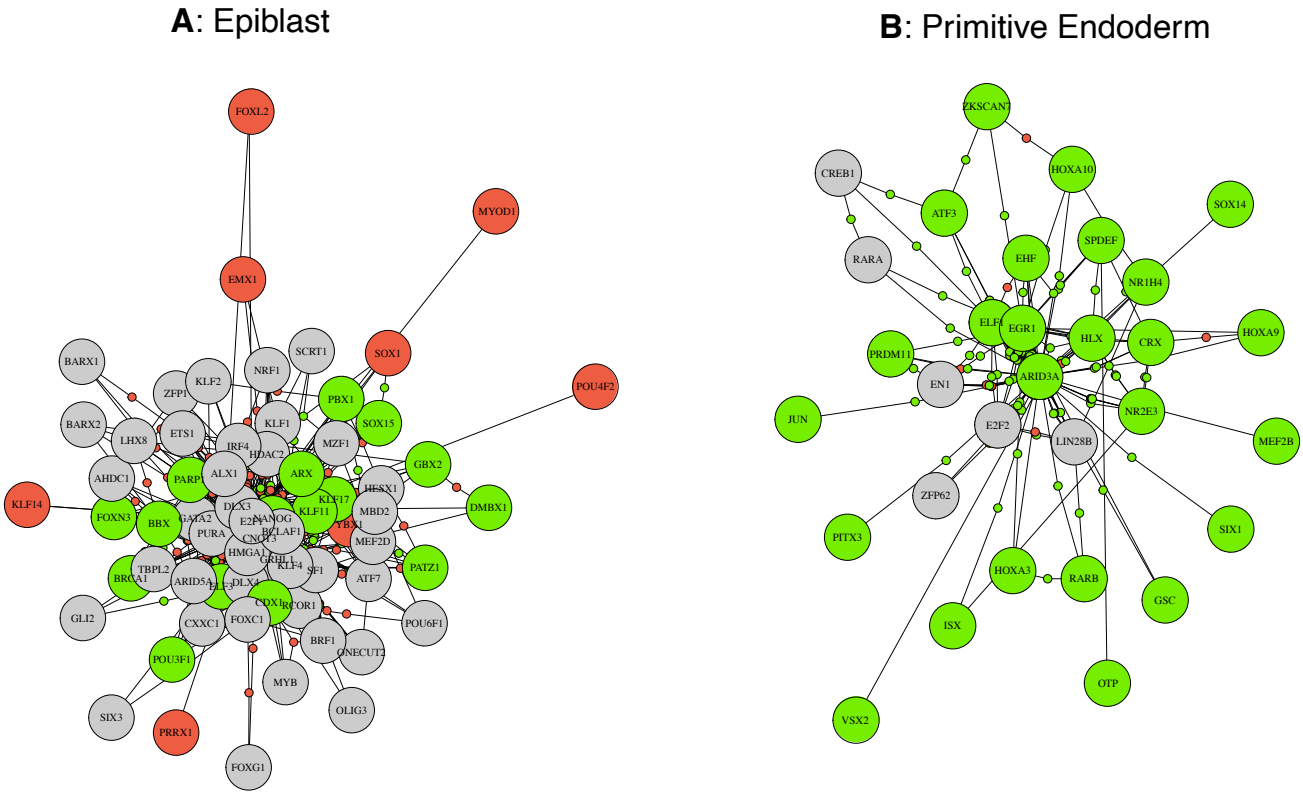

Figure 6: *(a) Epiblast (Epi) and (b) primitive endoderm (PE) differential regulation networks. Large circles indicate significantly differentially regulating transcription factors (Table S2); green and red indicate (respectively) significantly upregulated and downregulated TFs and genes; significance level is FDR* $p\text{-}val < 0.05$*, (t-test, Benjamini-Hochberg adjustment).*

2.5 Breast cancer at-risk prior epigenomic network

Luminal progenitor, mature luminal, and basal cells are epithelial cell subtypes that are thought to have important roles in breast cancer [47], and we have previously developed and validated DNAme cancer biomarkers based on definitions of these cell types [27, 34]. We inferred epigenomic prior networks for each of the cell types luminal progenitor, mature luminal, basal, and fibroblasts, using MethCorTarget with the breast cancer at-risk DNAme dataset [48], and using MLC with the breast cancer at-risk H3K27ac ChIP-seq dataset [48]. We chose to validate the DNAme epigenomic prior network against an epigenomic prior network inferred from H3K27ac ChIP-seq data, because H3K27ac is a histone modification that is associated with activated gene-regulatory regions of the DNA [49, 50]. We found a highly significant overlap between the TF-gene regulations inferred by MLC applied to the H3K27ac data and the regulations inferred by MethCorTarget applied to the DNAme data in all cell types, hence validating the inferred DNAme epigenomic prior networks. The odds-ratios for the overlap between the gene regulations inferred from H3K27ac data and DNAme data according to these methods are $\text{OR} = 45.5\ \ (44.8 - 46.2)$, $\text{OR} = 45.7\ \ (45.0 - 46.5)$, and $\text{OR} = 43.2\ \ (42.4 - 44.0)$, in luminal progenitor, mature luminal, and basal cells (respectively); statistical significance is $p < 0.001$ (Fisher's exact test) for all cell types. We also note that, interestingly, the epigenomic prior networks for luminal progenitor, mature luminal, and basal cells inferred from DNAme data again show candidate TF cis-regulations for at least eight time as many genes (respectively, 34792, 34680, and 34021 genes) than for chromatin accessibility (H3K27ac) data (respectively 4222, 4194, and 3597 different genes). Histograms showing the numbers of target-gene regulations for each TF, and the number of TF regulators for each target-gene, in the epigenomic prior networks for the breast cancer at-risk dataset, are shown in Figure Figure S5.

2.6 Breast cancer at-risk GRN inference and ablation study

Women who are carriers of a mutation in the *BRCA1* (*FANCS*) or *BRCA2* (*FANCD1* genes have a greatly increased chance of suffering from breast cancer during their lifetime [51]. Hence, apparently healthy cells from *BRCA1/2* mutation carriers are likely to harbour pre-cancerous genomic changes, and insights can be gained into how breast cancer develops by comparing healthy cells from *BRCA1/2* mutation carriers (BRCAmut) with cells from *BRCA1/2* wild-type (WT). Separately for BRCAmut and WT, using the cell type specific DNAme epigenomic prior networks obtained from the DNAme breast cancer at-risk dataset [48] for luminal progenitor, mature luminal, and basal epithelial cell subtypes, as well as fibroblast cells, we inferred GRNs as described (Section 4) from the combined breast cancer at-risk RNA-seq datasets [52–58], based on the definitions of these cell types shown in Figure S4. We note that the fibroblast cells that we have labelled '*type 1*' and '*type 2*' are likely to correspond to (respectively) interlobular and lobular fibroblasts, based on CD26 and CD105 expression (assessed here via respectively DPP4 and ENG mRNA expression in scRNA-seq data) [59]. Lobular fibroblasts share similarities with mesenchymal stem cells; they support epithelial growth and morphogenesis, and provide a specialised microenvironment for luminal epithelial progenitors [59]. Separately for each of these epithelial and fibroblast cell-subtypes, we inferred GRNs from scRNAseq data separately for *BRCA1/2* wild-type (WT) and *BRCA1/2* mutation carriers (BRCAmut).

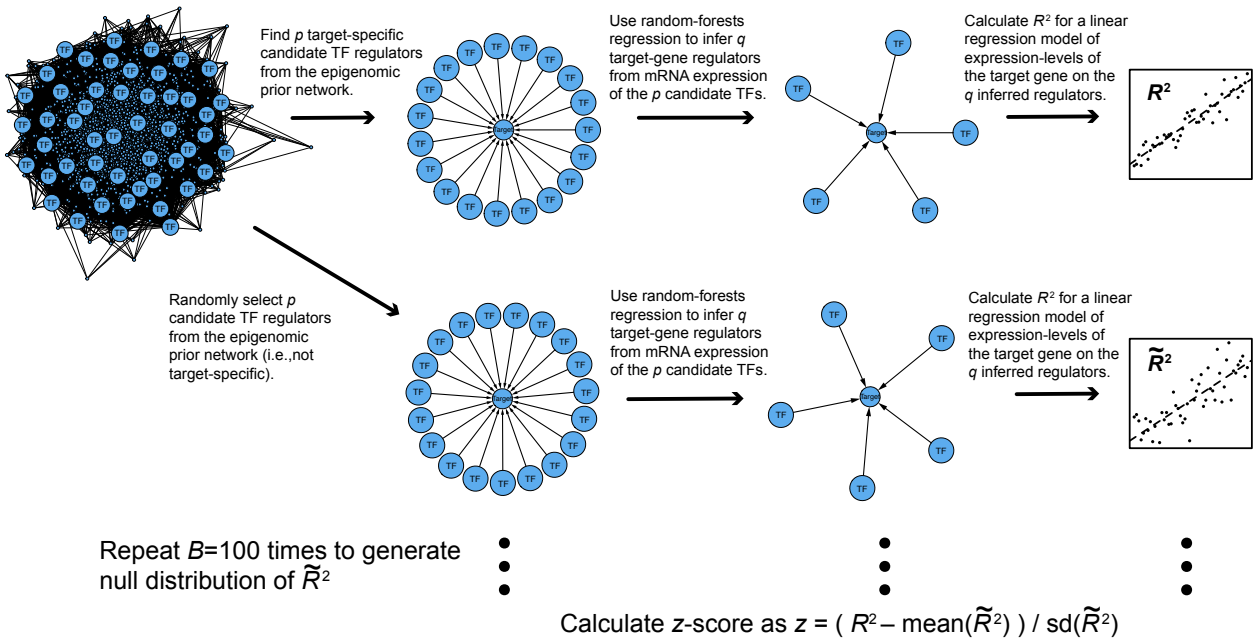

Figure 7: *Overview of the ablation study. As in our proposed methodology (Figure 1), separately for each target gene, the epigenomic prior network defines p candidate regulators. Random-forests regression is then used to choose $q < p$ of these p candidates, with the expression of these p TFs as predictors and the target-gene expression as the response. The R^2 statistic quantifies the proportion of variance in the target-gene expression, that is explained by the expression of these q inferred regulators, in a linear model. A null-distribution is then generated for comparison, by randomly selecting p predictors from all those found in the epigenomic prior network (i.e. not target-specific), then applying the rest of the procedure described above to calculate a null-distribution R^2 statistic in the same way (which we call $\tilde{R}^2$). This randomisation is repeated 100 times to obtain a null-distribution of $\tilde{R}^2$. A z-score for the R^2 statistic for this target-gene is then calculated using the mean and standard-deviation of these 100 null-distribution samples of $\tilde{R}^2$.*

We carried out an ablation study (Figure 7) using the combined breast cancer at-risk dataset, to test the claim that the prior epigenomic network can significantly inform gene-regulatory network inference. We calculated R^2 statistics to quantify how much of the variation in target-gene expression is explained by the TF regulators inferred for this gene target in our GRN. We then compared these observed R^2 against a null distribution of R^2 statistics which we call $\tilde{R}^2$, which are generated by applying our proposed GRN inference procedure to randomisations of the epigenomic prior network. We hence use the observed R^2 with the mean and standard-deviation of the null-distribution of $\tilde{R}^2$ statistics to calculate a z-score for each target-gene. T-tests show that genome-wide, these z-scores are significantly greater than 0 (Table 1), consistently across cell-types and conditions, demonstrating that the prior network significantly informs gene-regulatory network inference.

	t-statistic	*p*-value
luminal progenitor WT	20.7	<0.001
luminal progenitor BRCAmut	13.5	<0.001
luminal mature WT	21.1	<0.001
luminal mature BRCAmut	13.6	<0.001
basal WT	19.3	<0.001
basal BRCAmut	6.52	<0.001
fibroblast type1 WT	21.4	<0.001
fibroblast type1 BRCAmut	11.4	<0.001
fibroblast type2 WT	34.6	<0.001
fibroblast type2 BRCAmut	26.3	<0.001

Table 1: *Results of the ablation study. Significance-test p-values and t-statistics show that the z-scores calculated in the ablation study are significantly greater than 0 over all target-genes genome-wide, consistently across cell-types and conditions, demonstrating the added value of the epigenomic prior network to the GRN inference.*

2.7 Differential regulation in cells-of-origin of breast cancer

Significantly differentially regulating (diff-reg) TFs (Table S3, Table S4, Table S5, Table S6, Table S7) were identified by testing the differences in regulatory strength (measured by random-forests importance gain) in BRCAmut compared to WT, for all genes with physical evidence that cis-regulation is possible by that TF according to the epigenomic prior network. Luminal progenitor cells are hormone-receptor negative cells that are known to be the cell-of-origin of aggressive triple-negative breast cancer (TNBC) [47], and *BRCA1/2* mutation carriers are known to harbour pre-cancerous changes to in DNAme profiles [24, 26]. Hence, cis-regulatory regions of target-genes of significantly diff-reg TFs in *BRCA1/2* mutation carriers (BRCAmut) compared to wild-type (WT) in luminal progenitor cells (Figure 8a and b and Table S3) may harbour pre-cancerous alterations to DNAme profiles (methylomes). Therefore, we verified the significance of the *inferred diff-reg TFs* with respect to established knowledge, as follows. As an example, we considered the most strongly hyper-regulating TF in BRCAmut (compared to WT) in luminal progenitor cells (Table S3), the TF regulator *BACH1*; other examples can be found similarly. *BACH1* has recently been shown to be expressed more highly in TNBC [60]. Given the elevated risk of breast cancer in *BRCA1/2* mutation carriers, together with the knowledge that luminal progenitors are the cell-of-origin of TNBC, we would expect to see signs of increased regulation (hyper-regulation) by *BACH1* in luminal progenitors in *BRCA1/2* mutation carriers, as observed here. Corruption of the DNAme architecture [28] often precedes malignant transformation in cells by providing an 'epigenetically permissible environment', also known as an 'epigenetic field defect' [24, 26], and hence pre-cancerous alterations to the methylome can can serve as prognostic biomarkers [27]. Therefore, the GRN inference methodology proposed here could be used as a basis on which to develop DNAme biomarkers of cancer early-warning. Hence, we verified the significance of *gene-targets of the inferred diff-reg TFs* with respect to established knowledge, as follows. As an example, we looked in more detail at the gene found to be co-regulated by the largest number of diff-reg TFs in BRCAmut (compared to WT) in luminal progenitor cells, the gene *PRR7* (Figure 8); other examples can be found similarly. The *PRR7* gene is hypo-regulated in BRCAmut compared to WT, with inferred regulation by 5 and 10 diff-reg TFs, respectively (Figure 8). The *PRR7* promoter has three CpGs annotated on the Illumina 450K DNAme microarray, and there are publicly-available TCGA reference data of this type for breast cancer (BRCA) [61]. Therefore, we tested the prognostic significance of these three CpGs in the TCGA BRCA data in a survival analysis (Cox regression adjusted for confounding due to sample cell-type heterogeneity and clinical covariates), finding that the methylation level of two out of the three CpGs is significantly associated with patient survival outcome (Figure 8 and Table S8). As this gene is hypo-regulated in the at-risk population (*BRCA1/2* mutation carriers), we would expect that promoter hypermethylation at this gene would indicate poor prognosis, as observed here.

We also investigated how somatic mutations at TF binding sites could play a role in differential gene-regulation in this at-risk (*BRCA1/2* mut) population. We used the SComatic software [62] to infer somatic mutations in individual cells of the combined breast-cancer at-risk dataset. Then, we calculated the average somatic mutation rate over the TF binding sites of the diff-reg TFs, for all the gene-regulations inferred for these diff-reg TFs in the final GRN. As expected, the somatic mutation rates are higher in BRCAmut than WT for most cell types (Figure 9). However, interestingly, there does not seem to be a significantly elevated somatic mutation rate in BRCAmut (compared to WT) in luminal epithelial cells, which are known to be the cells-of-origin of aggressive breast cancers. This highlights the how important epigenomic changes are to gene regulation as a precursor to these cancers, in at-risk populations such as *BRCA1/2* mutation carriers [24, 26, 27].

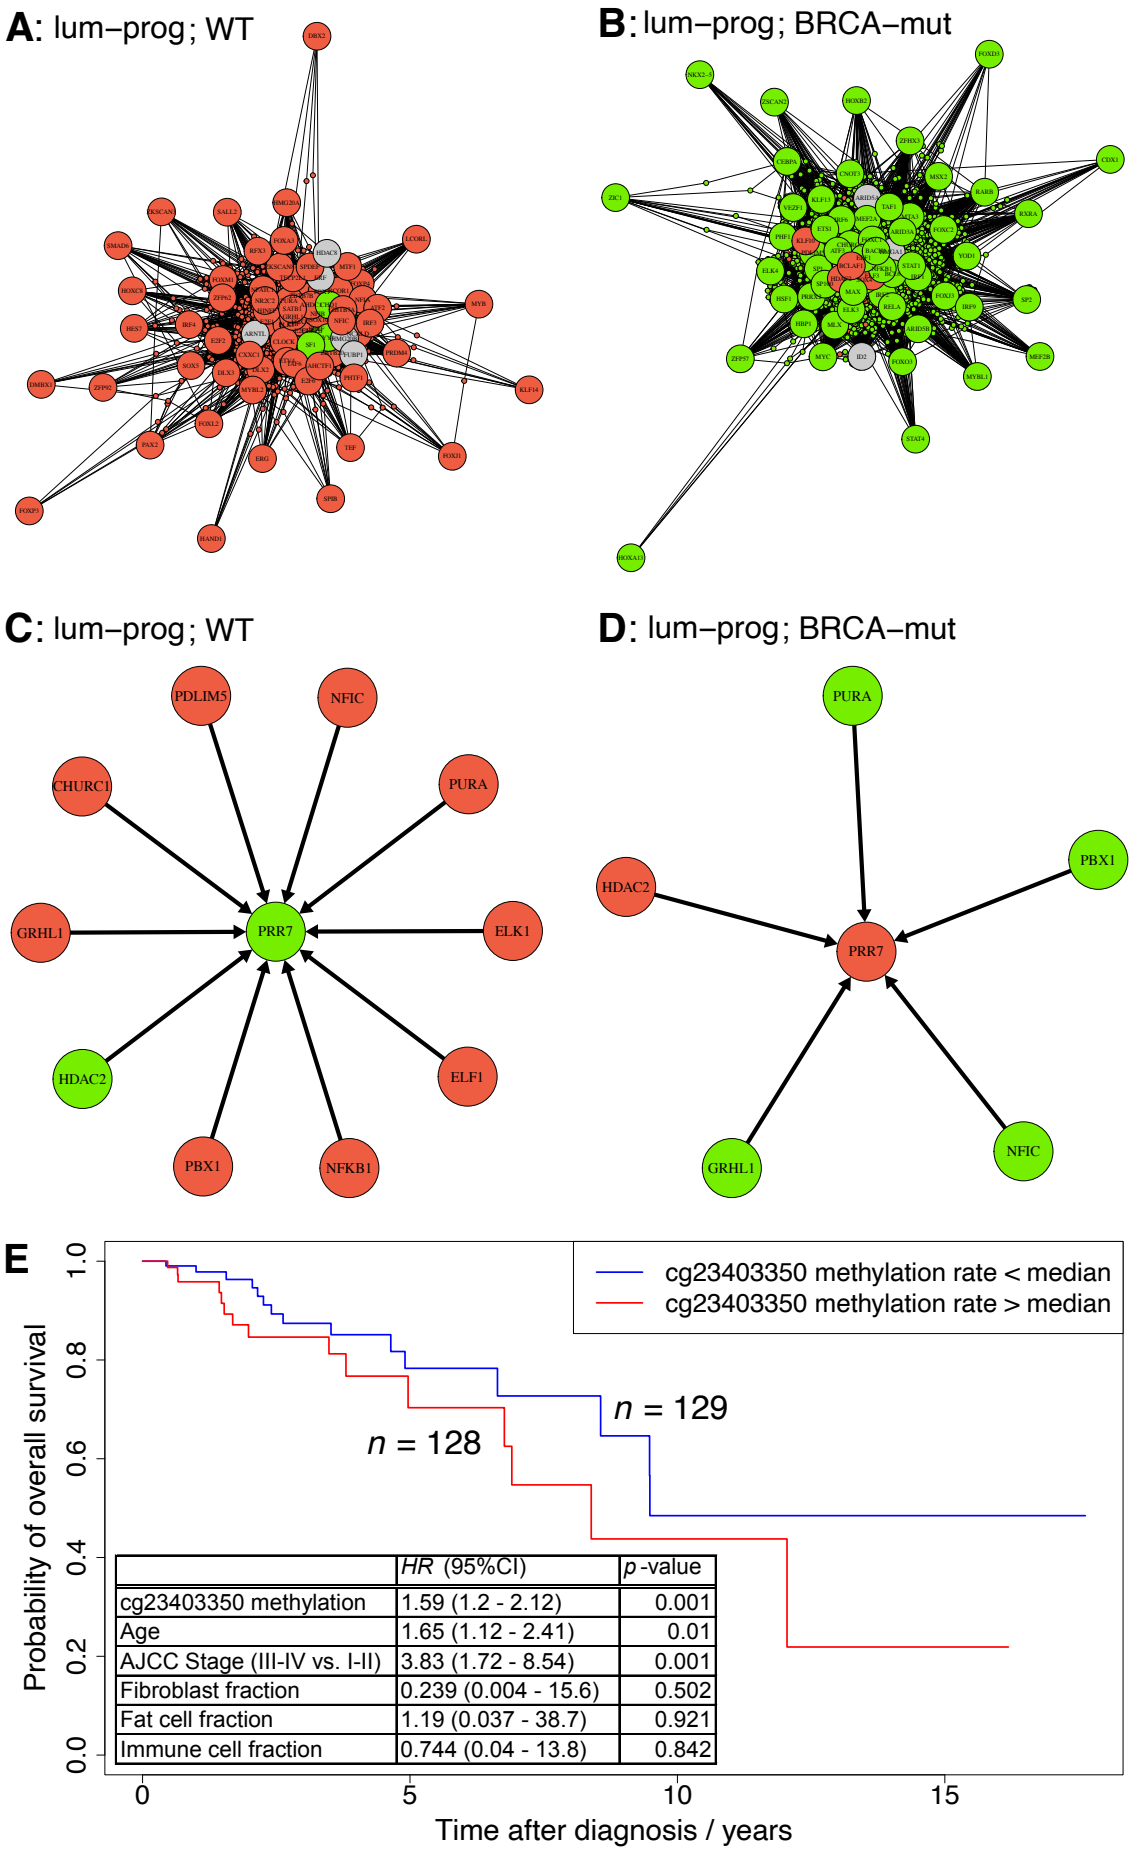

	HR (95%CI)	*p*-value
cg23403350 methylation	1.59 (1.2 - 2.12)	0.001
Age	1.65 (1.12 - 2.41)	0.01
AJCC Stage (III-IV vs. I-II)	3.83 (1.72 - 8.54)	0.001
Fibroblast fraction	0.239 (0.004 - 15.6)	0.502
Fat cell fraction	1.19 (0.037 - 38.7)	0.921
Immune cell fraction	0.744 (0.04 - 13.8)	0.842

Figure 8: *The gene-target PRR7 was found to have the largest number of regulations by differentially regulating (diff-reg) TFs of all genes, genome-wide. It is hence used as a case-study building on known, published results. (a) and (b) show the differential-regulation networks for luminal-progenitor cells for WT and BRCAmut (respectively). I.e., the part of the inferred GRNs containing differentially regulating (diff-reg) TFs (t-test, adj-p* < 0.05*) and differentially expressed target-genes (LIMMA, adj-p* < 0.05*). TFs and target-genes are shown with large and small circles respectively. (c) and (d) show the part of the GRN containing the target-gene PRR7 and its inferred diff-reg TF regulators, in WT and BRCAmut respectively. (e) Survival analysis (Cox regression adjusted for confounding by cell-type heterogeneity and clinical covariates) based on the TCGA BRCA dataset [61], for one of three CpGs (cg23403350) that are annotated to the PRR7 promoter on the Illumina 450K DNAme microarray (two of these three CpGs, cg23403350 and cg04110559, are significantly associated with survival outcome, Table S8).*

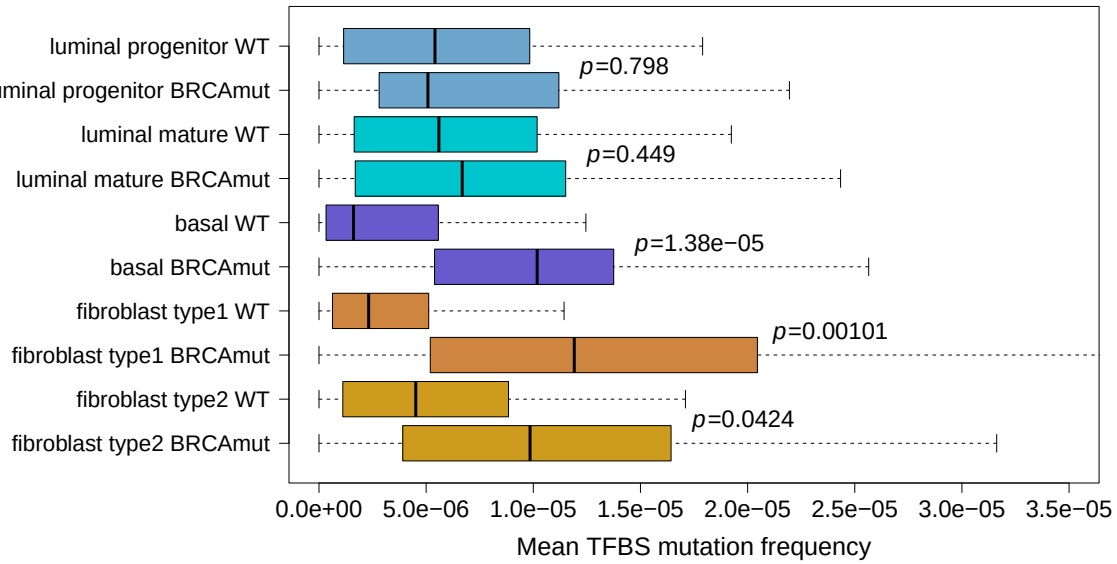

Figure 9: *Average mutation frequencies in specific cell-types, over the TF binding sites of the differentially regulating (diff-reg) TFs (t-test, adj-p* < 0.05*), for all the gene-regulations inferred for these diff-reg TFs in the final GRN.*

3 Discussion

In this work, we have shown that the accuracy of inference of GRN TF-gene cis-regulations is improved when epigenomic prior network structure is used in combination with advanced regression methods and gene-expression data. Furthermore, we have shown that DNAme data can be used successfully and effectively to estimate epigenomic prior network structure. There is a wealth of existing and publicly-available DNAme datasets for a wide range of tissue types that can now be leveraged using our proposed methodology to improve inference of gene-regulatory processes [13]. This contrasts with novel ATAC-seq datasets that are expensive to generate, and are currently less numerous in the public domain. However, as the availability increases of single-cell ATAC-seq data and of matched single-cell ATAC-seq and RNA-seq data, it will be interesting to see how our proposed methodology may be complemented and enhanced by the addition of single-cell chromatin accessibility data. For example, we note that single-cell ATAC-seq data are typically very sparse. Therefore, a large number of cis-regulatory regions would be expected to have low coverage in such data, and particularly in these regions the DNAme epigenomic prior network may provide useful complementary information.

We have also shown that with real data from well-studied contexts, the most significant findings from our methodology are well supported by existing literature. This gives confidence that our methodology will provide verifiable inferences in new contexts. However, we also emphasise that novel findings from computational GRN inferences remain hypotheses until they are experimentally validated. Further investigation of the differentially regulating transcription factors and their targets shown in Figure S8 and Table S3, Table S4, Table S5, Table S6, and Table S7, may allow gene-regulatory processes to be identified that are important to the oncogenic potential of these cell-types in *BRCA1* (*FANCS*) and *BRCA2* (*FANCD1*) mutation carriers, as well giving clues about possible pharmacological interventions. For example, Figure 8 shows how the methodology proposed here can be used to identify CpG loci that show altered methylation levels prior to cancer formation in at-risk tissues and populations. Hence, testing the DNAme rate of CpGs in at-risk individuals in certain cis-regulatory regions of tumour suppressor

genes or oncogenes, identified by the methodology proposed here, could provide important early-warning biomarkers of cancer or of cancer predisposition. However, further work will be required to systematically find and then validate all the examples of changes in DNAme patterns that may be discovered with our proposed methodology, and that are associated with changes in TF cis-regulation of specific genes that precede the onset of disease.

Previous advanced methods for GRN inference, such as GENIE3 [5], inferred GRNs based only on gene expression data. By including TF motif-DNA matches along with epigenomic data such as DNAme data, we make our GRN inferences more physically realistic. However, chromatin state (such as histone modifications), as well as the interaction of co-factors, and sequence context, all play an important part in determining gene regulation. Therefore, future development of our methodology will incorporate a wider range of 'omics and sequence data to further refine the GRN inferences. While we have shown that our proposed methodology is effective in the context of datasets of a wide range of sizes from the very different contexts of human embryonic development, healthy individuals at-risk of breast cancer, and individuals with advanced breast cancer, we acknowledge that this methodology remains untested outside of these contexts. Future work planned, includes expanding our proposed methodology in a tailored way to a wide range of other applications. This will include making a wide range of epigenomic prior networks available for different cell-types.

4 Methods

4.1 Quantification and statistical analysis

4.1.1 Epigenomic prior network construction

To obtain the candidate TF cis-regulators of a particular gene, we build on established methods called lever [22] and cistarget [23]. Those methods examine whether the association between TF binding motifs and cis-regulatory DNA sequences corresponds to the association between sample-specific observations and the same cis-regulatory DNA regions, such as open chromatin (as inferred from ATAC-seq or LiCAT-seq data). Lever [22] and cistarget [23] approach this as a prediction problem: i.e., does the matching of TF motif and DNA sequence at a particular location predict epigenomic patterns (such as open chromatin) that are observed at that location. In lever [22] and cistarget [23], the accuracy of this prediction is measured by the AUC (area under receiver-operator curve) statistic. The AUC statistic can be calculated from the area under the curve that plots the true-positive rate against the false-positive rate, as the classification threshold is decreased from infinity to zero; i.e., the receiver-operator curve. Equivalently, the AUC statistic can be calculated by a linear transformation of Somers' D statistic by the transformation $\text{AUC} = (D+1)/2$ [63], where $D \in [-1, 1]$. Somers' D statistic can also be thought of as equivalent to a correlation coefficient between a continuous variable (the prediction) and a binary variable (the ground-truth). If the correlation coefficient $\rho \in [-1, 1]$ is a measure of association between two binary variables, then $D \in [-1, 1]$ can be thought of as an analogous measure of association between a binary variable and a continuous variable. From the perspective of the AUC statistic, the binary variable corresponds to a representation of the true class $y^{(r)} \in \{0, 1\}$ for rankings $r \in \{1, ..., R\}$, and the continuous variable $x^{(r)}$ is used to predict that class, where a value of the AUC statistic close to 1 indicates that high rankings mostly correspond to a true class of 1.

In MethCorTarget, we use the absolute Spearman correlation coefficient to assess whether the average DNAme rate $y_j^{(r)} \in [0, 1]$ for sample $j \in \{1, ..., n\}$ at cis-regulatory region $r \in \{1, ..., R\}$ corresponds to the closeness of the match $x_l^{(r)}$ between TF motif $l \in \{1, ...L\}$ and the DNA sequence at that region, where correlation coefficient $\rho_{l,j}(x_l^{(r)}, y_j^{(r)})$ is calculated over $r \in \{1, ..., R\}$. In modified lever-cistarget (MLC), we equivalently use Somers' D to assess whether the presence of open chromatin (assessed by LiCAT-seq or H3K27ac) at a cis-regulatory region corresponds to the closeness of the match between the TF motif and the DNA sequence at that region. After calculating $\rho_{l,j}$ (for a given DNAme sample j) or $D_{l,j}$ (for other epigenomic data sample j) for TF motif l, we apply the Fisher transform to obtain approximately standard Gaussian-distributed variables $z_{l,j} \sim \mathcal{N}(0, 1)$ [64], $l \in \{1, ...L\}$, $j \in \{1, ..., n\}$, for L TFs and n samples (assuming the distributional assumptions also hold sufficiently well for D), and then calculate corresponding p-values from a standard Gaussian distribution. We then retain TFs in the subsequent analysis if the corresponding z-test leads to p-value $p_{l,j} < 0.1$, noting that this threshold provides optimum identifiability of network structures [65]. We also note that this p-value threshold might result in many false-positives for a large number of TFs, but that many of these would be expected to be filtered out in the next step in the GRN inference using gene expression data with advanced regression methods. The aim of the epigenomic prior network is to identify candidate regulators of a target-gene, which are then chosen from by the advanced regression method as it is applied to gene-expression data. We refer to this method of epigenomic prior network inference as MLC (modified-lever-cistarget) when the method calculates Somers' D from chromatin accessibility data (or similar), and we refer to this method as MethCorTarget when the method calculates the Spearman correlation coefficient ρ from DNAme data.

Histograms showing the numbers of target-gene regulations for each TF, and the number of TF regulators for each target-gene, in the epigenomic prior networks for the breast cancer at-risk dataset, and for the human embryonic development dataset, are shown in Figure Figure S5.

4.1.2 Gene-regulatory network inference

Gene-regulatory network (GRN) inference was carried out to infer TF-gene regulations by modelling the expression of each target-gene as a (possibly non-linear) combination of the expression of a set of candidate regulators (Equation 1) [4]. The target-gene approach to GRN inference [4] uses advanced regression methods such as random forests regression [66], and recently the xgboost implementation of random forests as gradient-boosted trees [38], to model the expression level of a particular 'target' gene (the response), conditional on the expression levels of a set of other genes (predictors, such as TFs), Equation 1. The target-gene approach to GRN inference is best known in the context of GENIE3 [5] and GRNboost2 [6]. In GENIE3 and GRNboost2, the set of candidate regulators corresponds to all the TFs in the dataset of interest, or to some other pre-specified list of candidate regulators that importantly is the same for all target-genes. In our proposed methodology, we use the epigenomic prior network to first infer a specific set of candidate TF regulators that can be different for each gene (Figure 1). This improves the accuracy of inference by the advanced regression method (Figure 3), because when using advanced regression such as random forests for variable selection, the advanced regression only needs to choose from 10s of candidate cis-regulating TFs. These 10s of candidate cis-regulators correspond to epigenomic prior information about possible interactions of the TF with the cis-regulatory DNA of the target-gene. This contrasts with the full list of 100s or 1000s of TFs that would need to be specified genome-wide when using GENIE3 [5] and GRNboost2 [6], as part of the SCENIC / SCENIC+ pipeline [1–3]. Finally, we construct the GRN by combining the local target-gene network model fits genome-wide, with each gene taking a turn as the response variable in the advanced regression, i.e., as the target-gene [4]. In the final GRN, we infer network edges using optimised thresholds on the random-forests importance-gain in each target-gene model fit. We estimated optimal thresholds based on a maximum-AUC criterion in data with known ground-truth network structure, for relevant values of sample size n and number of parents in the network p (Figure S6). The practitioner is free to choose the threshold as a user parameter, based on considerations of the downstream analyses. We choose to set the threshold as the lower-quartile of the optimal AUC value conditional on sample size n and number of parents in the network p. This lower choice of threshold ensures that false-negatives are reduced at the price of additional false-positives, which is preferred for two reasons. Firstly, further downstream validation will be used, mitigating any increase in false-positives. Secondly, downstream we analyse aggregations of gene regulations, which also makes the downstream analyses more tolerant to false-negatives, as their effect is mitigated when aggregated with true positives (Figure S6).

We found previously [12] that advanced regression methods as part of the target-gene approach to GRN inference perform best with normalisation methods such as FPKM and TPM. In this work, we chose to use CPM normalisation rather than FPKM or TPM normalisation, because inferences from advanced regression with the random forests method should not be affected by monotonic transformations of the data, such as scaling by gene length. In comparison with that previous work [12], we also chose to use random forests regression [66] (xgboost implementation [38]) rather than mutual information, for three reasons as follows: (1) The computational efficiency of xgboost [38]) makes it feasible to use xgboost with large, modern single-cell datasets such as the breast cancer at-risk dataset we analyse in this work. This contrasts with the much smaller number of cells in the human embryonic development dataset, where cell numbers are restricted due to the very precious and restricted nature of the source material; (2) In our previous work [12] we used the GENIE3 implementation [5] of target-gene GRN inference with random-forests [66], but in this work we have have shown a clear improvement over GENIE3 on synthetic data by using prior network structure (Figure 3), noting that GENIE3 does not use prior network structure that is tailored to each target-gene; (3) In this work we have based our epigenomic prior network inference methods on the lever [22] / cistarget [23] approach, which is also used together with random forests [66] as part of the well validated scenic [1, 3] pipeline. On the other hand our previous work [12] used RGT-HINT [36] to infer candidate TF-gene cis-regulations.

4.1.3 LiCAT-seq data processing

For the human embryonic development LiCAT-seq dataset [39], the DNA loci of cis-regulatory TF motifs corresponding to open chromatin were inferred as part of an earlier study [12] (with alignment to the GRCh38 reference human genome), and were reused in this work. These TF motif footprints at open chromatin were originally inferred from the human embryonic development LiCAT-seq dataset available from the NCBI Sequence Read Archive (SRA) under accession number SRP163205, from the URL `https://www.ncbi.nlm.nih.gov/sra/?term=SRP163205`
The TF-motif footprints at open chromatin were originally obtained using rgt-hint v0.13.0 [36] with TF binding models from HOCOMOCO and JASPAR [39], finding 433270 candidate cis-regulations identified over 13950 different genes by 894 different TFs.

We applied pycistarget and modified-lever-cistarget (MLC) to the downstream peak-calls (54425 peak-calls genome wide) for the human embryonic development LiCAT-seq dataset [39] for comparison. When applied to these peak-calls, pycistarget returned 124 significant TF motifs across 39298 distinct genomic regions. The MLC method found 549 significant TF motifs, matched across 27280 distinct genomic regions (z-test, $p < 0.1$, we note that many false positives generated in this step are expected to be removed by the subsequent GRN inference with gene-expression data). For MLC, candidate cis-regulations at these 27280 genomic regions were inferred from the

closest gene promoters ($<$ 5kb, according to Ensemble 107 [67]). This lead to a epigenomic prior network for the human embryonic development LiCAT-seq dataset, representing 21381 candidate cis-regulations of 2650 different genes by 270 different TFs.

4.1.4 H3K27ac ChIP-seq data processing

For the breast cancer at-risk H3K27ac ChIP-seq dataset [48], raw fastq files of unaligned sequencer reads for purified cell samples for the cell types luminal progenitor, mature luminal, and basal from each of 3 women (aged 22, 37 and 41), were downloaded from EGA (the European Genome-phenome Archive, `https://ega-archive.org`), under accession number EGAS00001000552. ChIP-seq reads were aligned to the GRCh38 (hg38) reference human genome assembly from Ensembl 109 [68] using the Burrows-Wheeler Alignment (BWA) software [69] (with default settings). Significant ChIP-seq peaks (FDR-adj $p < 0.05$) were then called using the callpeak function from the macs2 software [70, 71] (with default settings), resulting in three sets of peak calls per cell type. Applying MLC to these data for the cell types luminal progenitor, mature luminal, and basal, lead to (respectively) 583, 591, and 595 significant TF motifs matched across 25706, 27668, and 22570 genomic regions (z-test, $p < 0.1$, we note that many false positives generated in this step are expected to be removed by the subsequent GRN inference with gene-expression data). For each genomic region, candidate cis-regulations were inferred from the closest gene promoters ($<$ 5kb, according to Ensemble 107 [67]), also requiring that the candidate cis-regulation should be observed at this significance level in a majority of the samples for that cell type. This lead to a epigenomic prior network for each of the cell types luminal progenitor, mature luminal, and basal, representing (respectively) 142766, 142237, and 110513 candidate cis-regulations of 4222, 4194, and 3597 different genes, by 282, 285, and 277 different TFs.

4.1.5 DNAme bisulphite-seq data processing

For the *human embryonic development* DNAme dataset [40], CpG read-count data tables summarising bisulphite-seq libraries (aligned to the hg19 reference human genome) for 103 human ICM cell samples (inner cell mass, comprising epiblast and primitive endoderm cells, one library per cell sample), were downloaded from from the Gene Expression Omnibus (GEO) under accession number GSE81233, from the URL:
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE81233`

For the *breast cancer at-risk* DNAme dataset [48], bisulphite-seq DNAme data for purified cell samples donated by 3 women (aged 22, 37 and 41, one library per cell type for each woman) for luminal progenitor, mature luminal, basal, and stromal (fibroblast) cells, were downloaded from EGA (the European Genome-phenome Archive, `https://ega-archive.org`), under accession number EGAS00001000552. The downloaded raw sequencer reads were aligned to the GRCh37 (hg19) reference human genome assembly from Ensembl 104 [72], and methylated and unmethylated reads were counted at CpG sites, using the Bismark software [73] (with default settings).

To generate the epigenomic prior networks from these DNAme datasets, MethCorTarget was applied to the raw methylated and unmethylated bisulphite-seq counts, as follows. First, the methylated and unmethylated counts were aggregated over all candidate TF-motif binding regions that were mapped by at least 10 reads in total, leading to an average DNAme rate for each candidate TF-motif binding region, for each sample (i.e., bisulphite-seq library). The average methylation rates over the samples were then calculated for all TF-motif binding regions, for only those regions with data for at least 3 cells in the human embryonic development dataset, and for only those regions with data for the majority of samples in the breast cancer at-risk dataset. We note that DNAme data are typically quite noisy, but by averaging the methylation rates over TF binding regions, we mitigate some of the effect of this noise, as shown previously [74].

We used MethCorTarget to test the correlation of the ranking of the TF-motif DNA sequence matches for each binding region, with the average methylation rates over the binding regions calculated as above. We applied MethCorTarget in the same way for ICM cells of the human embryonic development dataset, and for the cell types luminal progenitor, mature luminal, basal, and stromal (fibroblast) cells of the breast cancer at-risk dataset. For ICM cells, MethCorTarget found 868 significant TF motifs matched across 95122 genomic regions. For luminal progenitor, mature luminal, basal, and stromal (fibroblast) cells, MethCorTarget found (respectively), 826, 780, 747, and 771 significant TF motifs matched across 206076, 201359, 195211, and 200102 genomic regions (z-test, $p < 0.1$; we note that many false positives generated in this step are expected to be removed by the subsequent GRN inference with gene-expression data). For each genomic region, candidate cis-regulations were inferred from the closest gene promoters ($<$ 5kb) to the motif-matched TFs (according to Ensemble 107 [67]). This lead to a epigenomic prior network for ICM representing 417320 candidate cis-regulations of 23575 genes by 375 TFs. It also lead to a epigenomic prior network for each of the cell types luminal progenitor, mature luminal, basal, and stromal (fibroblast) representing (respectively) 1035073, 946625, 920222, and 991998 candidate cis-regulations of 34792, 34680, 34021, and 34623 mRNA transcripts, by 413, 370, 382, and 387 TFs.

For plotting the bisulphite-seq reads aligned to hg19 on hg38 (GRCh38) co-ordinates, the liftover package in R was used to convert CpG loci from the hg19 coordinates to hg38. We note that comparisons between LiCAT-seq data aligned to hg38 co-ordinates and DNAme data aligned to hg19 co-ordinates were at the level of unique identifiers for cis-regulatory regions in the TF-motif DNA-binding database, mitigating alignment errors and loss of data that could arise during coordinate conversion.

4.1.6 RNA-seq data processing

The *human embryonic development* RNA-seq dataset [42] was downloaded from ArrayExpress under accession number EMTAB3929, from URL

`https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-MTAB-3929`

After quality control and filtering to retain cells from embryos at 5-7 days post-fertilisation, data for 20407 features (RNA transcripts) in 1258 samples (cells) were identified for use in downstream GRN inference.

The breast cancer at-risk RNA-seq datasets [52–58] have been analysed together in combination previously [57], and hence are referred to here as the *combined breast-cancer at-risk* dataset, and were downloaded from (respectively) the Gene Expression Omnibus (GEO) under accession numbers GSE161529, GSE174588, GSE195665, GSE198732, GSE180878, and ArrayExpress under accession numbers EMTAB13664, EMTAB9841, from (respectively) URLs

`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE161529`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE174588`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE195665`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE198732`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE180878`
`https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-MTAB-13664`
`https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-MTAB-9841`

After quality control and filtering to retain data from healthy pre-menopausal nulliparous women, data for 19995 features (RNA transcripts) in 70725, 65786, and 57906 samples (cells) from (respectively) 11 *BRCA1* (*FANCS*) mutation carriers, from 5 *BRCA2* (*FANCD1*) mutation carriers, and from 18 *BRCA1* (*FANCS*) and *BRCA2* (*FANCD1*) wild-type volunteers, were identified for use in downstream GRN inference.

Quality control was carried out as follows, to retain only high-quality samples (cells) and features (RNA transcripts). For both the human embryonic development dataset and the combined breast cancer at-risk dataset, samples (cells) were retained by filtering out any cells with non-zero counts for fewer than 750 RNA transcripts, filtering out cells for which the most highly expressed transcript excluding MALAT1 account for at least 10% of the total counts, filtering out cells for which at least 10% of the total counts correspond to transcripts encoding mitochondrial proteins, filtering out cells for which at least 50% of the total counts correspond to transcripts encoding ribosomal proteins, and filtering out doublets [75]. For the human embryonic development dataset, features were filtered out if they have non-zero counts in fewer than 10 samples (cells), noting the smaller sample size and greater read-depth for this dataset. For the combined breast-cancer at-risk dataset, RNA transcripts were filtered out if they have non-zero counts in fewer than 100 cells. For the combined breast cancer at-risk dataset, the COMBAT software [76] was used to control batch effects while combining datasets.

After quality control, differential expression analysis was carried out using the limma software [77] from the edgeR package [78] in R, with significance level set according to false discovery rate-adjusted $p < 0.05$. Separately after quality control, hypervariable mRNA transcripts were identified for dimensional reduction, visualisation, and clustering, following the principle of selecting genes with the greatest biological variability [79]. We use an estimator of the amount of biological variability compared to technical variability (i.e., overdispersion [79]) in each mRNA feature $i \in \{1, ..., p\}$, defined here as $\log V_i / \log m_i$ where V_i and m_i are (respectively) the empirical variance and mean for feature i. We note that there is currently no consensus on the best-practice for selecting the most informative features using models of overdispersion, and that the best solution is likely to depend both on the downstream processing and on as the datasets that are analysed [80–83]. The top 2000 mRNA features were retained in this way for the breast cancer at-risk dataset, and the top 5000 mRNA features were retained for the embryonic development dataset (recognising that more features are needed for estimating the Laplacian eigenspace as this dataset contains many fewer cell samples).

The Laplacian eigenspace (LE) is the basis for the visualisations and clustering shown in Figure S1, Figure S2, and Figure S4, and is estimated as the subspace of the data defined by the eigendecomposition of the graph-Laplacian representation of the datapoints [84], according to the eigenvectors that correspond to the eigenvalues that explain at least 1% of the total variance of the data, after removing the mean component. The single-cell data for the specific cell types analysed in the GRN inference are identified from the LE with unsupervised learning by fitting a Gaussian mixture model in the Laplacian eigenspace (GMM-LE clustering) [84, 85], as shown in UMAP projections [86] from the LE [84] that appear in Figure S1, Figure S2, and Figure S4.

For the human embryonic development dataset, the epiblast cells used in the GRN inference are the same ones that we used in previous work [84], identified as cluster 5 in Figure S1; the identifiers for these Epi cells from the original study that generated those data [42] are listed in Table S9. GMM-LE clustering was then applied hierarchically to cluster 2 from Figure S1, resulting in 37 primitive endoderm (PE) cells identified as cluster 4 in Figure S2; the identifiers for these PE cells from the original study that generated these data [42] are listed in Table S10. The 68 and 37 Epi and PE cells (respectively) are validated in Figure S3 by the expression levels of known Epi marker genes *NANOG* and *KLF17* and PE marker genes *GATA4* and *SOX17* [12], and are used as representative samples of gene-expression profiles of those cell types in the GRN inference. For the breast cancer at-risk dataset,

we previously identified 5531 luminal progenitor, 2535 mature luminal, and 3913 basal cells [84], which are identified as clusters 1, 2, and 8, respectively in Figure S4. These cell samples are then used as representative examples of gene-expression profiles of those cell types in the GRN inference.

4.1.7 DNAme bulk microarray data processing

Bulk-tissue Illumina 450K DNAme microarray data for 257 breast cancer invasive carcinoma (BRCA) samples with matched clinical data were downloaded from the TCGA repository (The Cancer Genome Atlas) [61], and were preprocessed as follows. Raw DNAme idat data files were loaded in R and background-corrected and normalised with BMIQ [87]. Probes were removed if they had $< 95\%$ coverage across samples, and any remaining probes with detection p-value > 0.05 were replaced by k-NN imputation, with $k = 5$.

4.1.8 TF-motif databases

The rankings for the matches of TF-motif to DNA sequence location for human, that we used with pycistarget, MLC (modified-lever-cistarget) and MethCorTarget, are based on 32765 unique TF motifs that were originally collated from 29 different motif databases as part of the SCENIC+ package [3], and were downloaded under the filenames
`hg38_screen_v10_clust.regions_vs_motifs.rankings.feather` (for the motif rankings) and
`motifs-v10nr_clust-nr.mgi-m0.001-o0.0.tbl` (for the TF-motif matches), from the URL:
`https://resources.aertslab.org/cistarget/`

4.1.9 Synthetic data generation

Synthetic gene expression data was generated, corresponding to known ground-truth GRN structure. These data were generated using the 'gene netweaver' software [88], which uses realistic differential-equations models to simulate the data according to models of chemical reaction kinetics; this software was used previously to generate data for the DREAM4 and DREAM5 network inference challenges [37] that were originally won by GENIE3 [5]. Using gene-net weaver, we generated 300 synthetic gene-expression profiles based on known GRN structure. These 300 synthetic gene-expression profiles were then analysed equivalently to 300 normalised single-cell libraries sequenced at high depth. The number of synthetic data samples, the maximum synthetic network size, and the number of candidate regulators of each target gene, were determined by the constraints of the data generating process

4.1.10 Deconvolution of bulk DNA methylation data

Sample fractions of epithelial cells, fibroblasts, immune cells and fat cells were estimated for the TCGA BRCA bulk DNAme microarray data using the 'Epidish' package in R [33]. We note that only the fractions of fibroblasts, immune cells and fat cells were used to adjust for sample cell-type fraction due to built-in redundancy in these estimates (i.e., the epithelial, fibroblast, immune, and fat fractions sum to 1 by definition).

4.1.11 Software

Downstream analyses were carried out in R (versions 4.1-4.3), except for the application of pycistarget to the LiCAT-seq data, which used Python version 3.10 and pycistarget version 1.0.

4.2 Key resources

The human embryonic development DNAme dataset is available from the Gene Expression Omnibus (GEO) under accession number GSE81233, from the URL `https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE81233`

The human embryonic development LiCAT-seq dataset is available from the NCBI Sequence Read Archive (SRA) under accession number SRP163205, from the URL `https://www.ncbi.nlm.nih.gov/sra/?term=SRP163205`

The breast cancer at-risk DNAme (bisulphite-seq), and H3K27ac (ChIP-seq) datasets, are available from EGA under accession number EGAS00001000552, from the URL `https://ega-archive.org`

The human embryonic development single-cell RNA-seq dataset is available from ArrayExpress under accession number E-MTAB-3929, from the URL `https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-MTAB-3929`

The breast cancer at-risk single-cell RNA-seq datasets are available from the Gene Expression Omnibus (GEO) under accession numbers GSE161529, GSE174588, GSE195665, GSE198732, GSE180878, and ArrayExpress under accession numbers EMTAB13664, EMTAB9841, from URLs
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE161529`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE174588`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE195665`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE198732`
`https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE180878`
`https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-MTAB-13664`
`https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-MTAB-9841`

The breast cancer invasive carcinoma (BRCA) DNAme dataset and associated demographic information is available from TCGA [61] from the URL `http://www.cbioportal.org/study/summary?id=brca_tcga`

The TF-motif DNA-sequence match rankings are available from the URL: https://resources.aertslab.org/cistarget/

R scripts implementing the epigenomic prior network construction, GRN inference, and differential cis-regulatory analysis, with required data-files for the human embryonic development datasets are available from https://zenodo.org/records/17253602, and implemented with required data-files for the breast cancer at-risk datasets from https://zenodo.org/records/17255982

5 Additional information

Author contributions

TB conceived of and designed the study. TB wrote the manuscript, with advice from QH. Statistical analyses and model development were carried out by TB, with contributions from ML, CS, YG, and QH. All co-authors approved this version of this manuscript.

Funding

Some of the work of TB during this project was funded by the MRC grant MR/P014070/1. The work of QH was supported by the Francis Crick Institute, which receives its core funding from Cancer Research UK (CC2074), the UK Medical Research Council (CC2074), and the Wellcome Trust (CC2074). The funding bodies had no role in the design of the study, collection, analysis, and interpretation of data, or writing of the manuscript.

Acknowledgements

The results published here are in whole or part based upon data generated by The Canadian Epigenetics, Epigenomics, Environment and Health Research Consortium (CEEHRC) initiative funded by the Canadian Institutes of Health Research (CIHR), Genome BC, and Genome Quebec. Information about CEEHRC and the participating investigators and institutions can be found at http://www.cihr-irsc.gc.ca/e/43734.html

The authors acknowledge the use of the UCL High Performance Computing Facilities, and associated support services, in the completion of this work.

Competing interests

The authors declare that there are no competing interests.

License

This manuscript is available under a Creative Commons CC BY-NC License, as described at https://creativecommons.org/licenses/by-nc/4.0/

6 References

[1] Aibar S, González-Blas CB, Moerman T, Huynh-Thu VA, Imrichova H, Hulselmans G, et al. SCENIC: single-cell regulatory network inference and clustering. Nature methods. 2017;14(11):1083-6.

[2] Van de Sande B, Flerin C, Davie K, De Waegeneer M, Hulselmans G, Aibar S, et al. A scalable SCENIC workflow for single-cell gene regulatory network analysis. Nature Protocols. 2020;15(7):2247-76.

[3] Bravo González-Blas C, De Winter S, Hulselmans G, Hecker N, Matetovici I, Christiaens V, et al. SCENIC+: single-cell multiomic inference of enhancers and gene regulatory networks. Nature Methods. 2023;20(9):1355-67.

[4] Dobra A, Hans C, Jones B, Nevins JR, Yao G, West M. Sparse graphical models for exploring gene expression data. Journal of Multivariate Analysis. 2004;90(1):196-212.

[5] Huynh-Thu VA, Irrthum A, Wehenkel L, Geurts P. Inferring regulatory networks from expression data using tree-based methods. PloS one. 2010;5(9):e12776.

[6] Moerman T, Aibar Santos S, Bravo González-Blas C, Simm J, Moreau Y, Aerts J, et al. GRNBoost2 and Arboreto: efficient and scalable inference of gene regulatory networks. Bioinformatics. 2019;35(12):2159-61.

[7] Brandman O, Meyer T. Feedback loops shape cellular signals in space and time. Science. 2008;322(5900):390-5.

[8] Bartlett TE, Olhede SC, Zaikin A. A DNA methylation network interaction measure, and detection of network oncomarkers. PloS one. 2014;9(1):e84573.

[9] Bartlett TE, Zaikin A, et al. Detection of epigenomic network community oncomarkers. The Annals of Applied Statistics. 2016;10(3):1373-96.

[10] Bartlett TE, Kosmidis I, Silva R. Two-way sparsity for time-varying networks with applications in genomics. The Annals of Applied Statistics. 2021;15(2):856-79.

[11] Bartlett T. Fusion of single-cell transcriptome and DNA-binding data, for genomic network inference in cortical development. BMC bioinformatics. 2021;22(1):1-19.

[12] Alanis-Lobato G, Bartlett TE, Huang Q, Simon C, McCarthy A, Elder K, et al. MICA: A multi-omics method to predict gene regulatory networks in early human embryos. Life Science Alliance. 2023;7(1):e202302415.

[13] Loyfer N, Magenheim J, Peretz A, Cann G, Bredno J, Klochendler A, et al. A DNA methylation atlas of normal human cell types. Nature. 2023;613(7943):355-64.

[14] Butte AJ, Kohane IS. Mutual information relevance networks: functional genomic clustering using pairwise entropy measurements. In: Biocomputing 2000. World Scientific; 1999. p. 418-29.

[15] Daub CO, Steuer R, Selbig J, Kloska S. Estimating mutual information using B-spline functions–an improved similarity measure for analysing gene expression data. BMC bioinformatics. 2004;5(1):1-12.

[16] Margolin AA, Nemenman I, Basso K, Wiggins C, Stolovitzky G, Dalla Favera R, et al. ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. BMC bioinformatics. 2006;7(1):1-15.

[17] Faith JJ, Hayete B, Thaden JT, Mogno I, Wierzbowski J, Cottarel G, et al. Large-scale mapping and validation of Escherichia coli transcriptional regulation from a compendium of expression profiles. PLoS biology. 2007;5(1):e8.

[18] Chan TE, Stumpf MP, Babtie AC. Gene regulatory network inference from single-cell data using multivariate information measures. Cell systems. 2017;5(3):251-67.

[19] Stuart JM, Segal E, Koller D, Kim SK. A gene-coexpression network for global discovery of conserved genetic modules. science. 2003;302(5643):249-55.

[20] Bartlett TE, Müller S, Diaz A. Single-cell co-expression subnetwork analysis. Scientific reports. 2017;7(1):15066.

[21] Buenrostro JD, Giresi PG, Zaba LC, Chang HY, Greenleaf WJ. Transposition of native chromatin for fast and sensitive epigenomic profiling of open chromatin, DNA-binding proteins and nucleosome position. Nature methods. 2013;10(12):1213-8.

[22] Warner JB, Philippakis AA, Jaeger SA, He FS, Lin J, Bulyk ML. Systematic identification of mammalian regulatory motifs' target genes and functions. Nature methods. 2008;5(4):347-53.

[23] Herrmann C, Van de Sande B, Potier D, Aerts S. i-cisTarget: an i ntegrative genomics method for the prediction of regulatory features and cis-regulatory modules. Nucleic acids research. 2012;40(15):e114-4.

[24] Bernstein C, Nfonsam V, Prasad AR, Bernstein H. Epigenetic field defects in progression to cancer. World journal of gastrointestinal oncology. 2013;5(3):43.

[25] Bartlett TE, Chindera K, McDermott J, Breeze CE, Cooke WR, Jones A, et al. Epigenetic reprogramming of fallopian tube fimbriae in BRCA mutation carriers defines early ovarian cancer evolution. Nature communications. 2016;7(1):1-10.

[26] Teschendorff AE, Gao Y, Jones A, Ruebner M, Beckmann MW, Wachter DL, et al. DNA methylation outliers in normal breast tissue identify field defects that are enriched in cancer. Nature communications. 2016;7(1):1-12.

[27] Bartlett TE, Evans I, Jones A, Barrett JE, Haran S, Reisel D, et al. Antiprogestins reduce epigenetic field cancerization in breast tissue of young healthy women. Genome medicine. 2022;14(1):1-18.

[28] Bartlett TE, Zaikin A, Olhede SC, West J, Teschendorff AE, Widschwendter M. Corruption of the intra-gene DNA methylation architecture is a hallmark of cancer. PloS one. 2013;8(7):e68285.

[29] Bartlett TE, Jones A, Goode EL, Fridley BL, Cunningham JM, Berns EM, et al. Intra-gene DNA methylation variability is a clinically independent prognostic marker in women's cancers. PloS one. 2015;10(12):e0143178.

[30] Gosden RG, Feinberg AP. Genetics and epigenetics—nature's pen-and-pencil set. Mass Medical Soc; 2007.

[31] Kaluscha S, Domcke S, Wirbelauer C, Stadler MB, Durdu S, Burger L, et al. Evidence that direct inhibition of transcription factor binding is the prevailing mode of gene and repeat repression by DNA methylation. Nature genetics. 2022;54(12):1895-906.

[32] Feinberg AP, Ohlsson R, Henikoff S. The epigenetic progenitor origin of human cancer. Nature reviews genetics. 2006;7(1):21-33.

[33] Teschendorff AE, Breeze CE, Zheng SC, Beck S. A comparison of reference-based algorithms for correcting cell-type heterogeneity in Epigenome-Wide Association Studies. BMC bioinformatics. 2017;18(1):105.

[34] Bartlett TE, Jia P, Chandna S, Roy S. Inference of tissue relative proportions of the breast epithelial cell types luminal progenitor, basal, and luminal mature. Scientific reports. 2021;11(1):1-10.

[35] Ogris C, Hu Y, Arloth J, Müller NS. Versatile knowledge guided network inference method for prioritizing key regulatory factors in multi-omics data. Scientific Reports. 2021;11(1):6806.

[36] Li Z, Schulz MH, Look T, Begemann M, Zenke M, Costa IG. Identification of transcription factor binding sites using ATAC-seq. Genome biology. 2019;20(1):1-21.

[37] Meyer P, Saez-Rodriguez J. Advances in systems biology modeling: 10 years of crowdsourcing DREAM challenges. Cell Systems. 2021;12(6):636-53.

[38] Chen T, Guestrin C. Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining; 2016. p. 785-94.

[39] Liu L, Leng L, Liu C, Lu C, Yuan Y, Wu L, et al. An integrated chromatin accessibility and transcriptome landscape of human pre-implantation embryos. Nature communications. 2019;10(1):364.

[40] Zhu P, Guo H, Ren Y, Hou Y, Dong J, Li R, et al. Single-cell DNA methylome sequencing of human preimplantation embryos. Nature genetics. 2018;50(1):12-9.

[41] Blakeley P, Fogarty NM, Del Valle I, Wamaitha SE, Hu TX, Elder K, et al. Defining the three cell lineages of the human blastocyst by single-cell RNA-seq. Development. 2015;142(18):3151-65.

[42] Petropoulos S, Edsgärd D, Reinius B, Deng Q, Panula SP, Codeluppi S, et al. Single-cell RNA-seq reveals lineage and X chromosome dynamics in human preimplantation embryos. Cell. 2016;165(4):1012-26.

[43] Allègre N, Chauveau S, Dennis C, Renaud Y, Meistermann D, Estrella LV, et al. NANOG initiates epiblast fate through the coordination of pluripotency genes expression. Nature Communications. 2022;13(1):3550.

[44] Aksoy I, Jauch R, Chen J, Dyla M, Divakar U, Bogu GK, et al. Oct4 switches partnering from Sox2 to Sox17 to reinterpret the enhancer code and specify endoderm. The EMBO journal. 2013;32(7):938-53.

[45] Ben-Porath I, Thomson MW, Carey VJ, Ge R, Bell GW, Regev A, et al. An embryonic stem cell–like gene expression signature in poorly differentiated aggressive human tumors. Nature Genetics. 2008;40(5):499-507.

[46] Guo G, Stirparo GG, Strawbridge SE, Spindlow D, Yang J, Clarke J, et al. Human naive epiblast cells possess unrestricted lineage potential. Cell stem cell. 2021;28(6):1040-56.

[47] Tharmapalan P, Mahendralingam M, Berman HK, Khokha R. Mammary stem cells and progenitors: targeting the roots of breast cancer for prevention. The EMBO journal. 2019;38(14):e100852.

[48] Pellacani D, Bilenky M, Kannan N, Heravi-Moussavi A, Knapp DJ, Gakkhar S, et al. Analysis of normal human mammary epigenomes reveals cell-specific active enhancer states and associated transcription factor networks. Cell reports. 2016;17(8):2060-74.

[49] Creyghton MP, Cheng AW, Welstead GG, Kooistra T, Carey BW, Steine EJ, et al. Histone H3K27ac separates active from poised enhancers and predicts developmental state. Proceedings of the National Academy of Sciences. 2010;107(50):21931-6.

[50] Bogdanović O, Fernandez-Miñán A, Tena JJ, de la Calle-Mustienes E, Hidalgo C, van Kruysbergen I, et al. Dynamics of enhancer chromatin signatures mark the transition from pluripotency to cell specification during embryogenesis. Genome research. 2012;22(10):2043-53.

[51] Kuchenbaecker KB, McGuffog L, Barrowdale D, Lee A, Soucy P, Dennis J, et al. Evaluation of polygenic risk scores for breast and ovarian cancer risk prediction in BRCA1 and BRCA2 mutation carriers. JNCI: Journal of the National Cancer Institute. 2017;109(7):djw302.

[52] Pal B, Chen Y, Vaillant F, Capaldo BD, Joyce R, Song X, et al. A single-cell RNA expression atlas of normal, preneoplastic and tumorigenic states in the human breast. The EMBO journal. 2021;40(11):e107333.

[53] Nee K, Ma D, Nguyen QH, Pein M, Pervolarakis N, Insua-Rodríguez J, et al. Preneoplastic stromal cells promote BRCA1-mediated breast tumorigenesis. Nature genetics. 2023;55(4):595-606.

[54] Kumar T, Nee K, Wei R, He S, Nguyen QH, Bai S, et al. A spatially resolved single-cell genomic atlas of the adult human breast. Nature. 2023;620(7972):181-91.

[55] Murrow LM, Weber RJ, Caruso JA, McGinnis CS, Phong K, Gascard P, et al. Mapping hormone-regulated cell-cell interaction networks in the human breast at single-cell resolution. Cell systems. 2022;13(8):644-64.

[56] Gray GK, Li CMC, Rosenbluth JM, Selfors LM, Girnius N, Lin JR, et al. A human breast atlas integrating single-cell proteomics and transcriptomics. Developmental cell. 2022;57(11):1400-20.

[57] Reed AD, Pensa S, Steif A, Stenning J, Kunz DJ, Porter LJ, et al. A single-cell atlas enables mapping of homeostatic cellular shifts in the adult human breast. Nature genetics. 2024;56(4):652-62.

[58] Twigger AJ, Engelbrecht LK, Bach K, Schultz-Pernice I, Pensa S, Stenning J, et al. Transcriptional changes in the mammary gland during lactation revealed by single cell sequencing of cells from human milk. Nature communications. 2022;13(1):562.

[59] Morsing M, Klitgaard MC, Jafari A, Villadsen R, Kassem M, Petersen OW, et al. Evidence of two distinct functionally specialized fibroblast lineages in breast stroma. Breast Cancer Research. 2016;18(1):108.

[60] Lee J, Yesilkanal AE, Wynne JP, Frankenberger C, Liu J, Yan J, et al. Effective breast cancer combination therapy targeting BACH1 and mitochondrial metabolism. Nature. 2019;568(7751):254-8.

[61] Network CGA, et al. Comprehensive molecular portraits of human breast tumours. Nature. 2012;490(7418):61-70.

[62] Muyas F, Sauer CM, Valle-Inclán JE, Li R, Rahbari R, Mitchell TJ, et al. De novo detection of somatic mutations in high-throughput single-cell profiling data sets. Nature biotechnology. 2024;42(5):758-67.

[63] Newson R. Parameters behind "nonparametric" statistics: Kendall's tau, Somers' D and median differences. The Stata Journal. 2002;2(1):45-64.

[64] Fisher RA. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. Biometrika. 1915:507-21.

[65] Gallagher I, Jones A, Bertiger A, Priebe CE, Rubin-Delanchy P. Spectral embedding of weighted graphs. Journal of the American Statistical Association. 2023:1-10.

[66] Breiman L. Random forests. Machine learning. 2001;45:5-32.

[67] Cunningham F, Allen JE, Allen J, Alvarez-Jarreta J, Amode MR, Armean IM, et al. Ensembl 2022. Nucleic acids research. 2022;50(D1):D988-95.

[68] Martin FJ, Amode MR, Aneja A, Austine-Orimoloye O, Azov AG, Barnes I, et al. Ensembl 2023. Nucleic acids research. 2023;51(D1):D933-41.

[69] Li H, Durbin R. Fast and accurate short read alignment with Burrows–Wheeler transform. bioinformatics. 2009;25(14):1754-60.

[70] Zhang Y, Liu T, Meyer CA, Eeckhoute J, Johnson DS, Bernstein BE, et al. Model-based analysis of ChIP-Seq (MACS). Genome biology. 2008;9(9):1-9.

[71] Gaspar JM. Improved peak-calling with MACS2. BioRxiv. 2018:496521.

[72] Howe KL, Achuthan P, Allen J, Allen J, Alvarez-Jarreta J, Amode MR, et al. Ensembl 2021. Nucleic acids research. 2021;49(D1):D884-91.

[73] Krueger F, Andrews SR. Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications. bioinformatics. 2011;27(11):1571-2.

[74] Wittenberger T, Sleigh S, Reisel D, Zikan M, Wahl B, Alunni-Fabbroni M, et al. DNA methylation markers for early detection of women's cancer: promise and challenges. Epigenomics. 2014;6(3):311-27.

[75] Germain PL, Lun A, Meixide CG, Macnair W, Robinson MD. Doublet identification in single-cell sequencing data using scDblFinder. f1000research. 2022;10:979.

[76] Johnson WE, Li C, Rabinovic A. Adjusting batch effects in microarray expression data using empirical Bayes methods. Biostatistics. 2007;8(1):118-27.

[77] Smyth GK. Linear models and empirical bayes methods for assessing differential expression in microarray experiments. Statistical applications in genetics and molecular biology. 2004;3(1).

[78] Robinson MD, McCarthy DJ, Smyth GK. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. bioinformatics. 2010;26(1):139-40.

[79] McCarthy DJ, Chen Y, Smyth GK. Differential expression analysis of multifactor RNA-Seq experiments with respect to biological variation. Nucleic acids research. 2012;40(10):4288-97.

[80] Townes FW, Hicks SC, Aryee MJ, Irizarry RA. Feature selection and dimension reduction for single-cell RNA-Seq based on a multinomial model. Genome biology. 2019;20(1):1-16.

[81] Hafemeister C, Satija R. Normalization and variance stabilization of single-cell RNA-seq data using regularized negative binomial regression. Genome biology. 2019;20(1):296.

[82] Lause J, Berens P, Kobak D. Analytic Pearson residuals for normalization of single-cell RNA-seq UMI data. Genome biology. 2021;22(1):22.

[83] Choudhary S, Satija R. Comparison and evaluation of statistical error models for scRNA-seq. Genome biology. 2022;23(1):27.

[84] Bartlett TE, Chandna S, Roy S. Stochastic networks theory to model single-cell genomic count data. arXiv preprint arXiv:230302498. 2023.

[85] Rubin-Delanchy P, Cape J, Tang M, Priebe CE, et al. A statistical interpretation of spectral embedding: the generalised random dot product graph. Journal of the Royal Statistical Society Series B. 2022;84(4):1446-73.

[86] McInnes L, Healy J, Melville J. Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:180203426. 2018.

[87] Teschendorff AE, Marabita F, Lechner M, Bartlett T, Tegner J, Gomez-Cabrero D, et al. A beta-mixture quantile normalization method for correcting probe design bias in Illumina Infinium 450 k DNA methylation data. Bioinformatics. 2013;29(2):189-96.

[88] Schaffter T, Marbach D, Floreano D. GeneNetWeaver: in silico benchmark generation and performance profiling of network inference methods. Bioinformatics. 2011;27(16):2263-70.