

Detection of cooperatively bound transcription factor pairs using ChIP-seq peak intensities and expectation maximization

Vishaka Datta^{*}, Rahul Siddharthan[†], and Sandeep Krishna^{*}

^{*}Simons Centre for the Study of Living Machines, National Centre for Biological Sciences, TIFR, Bengaluru 560065, India

[†]The Institute of Mathematical Sciences/HBNI, Taramani, Chennai 600 113, India

May 18, 2017

Abstract

Transcription factors (TFs) often work cooperatively, where the binding of one TF to DNA enhances the binding affinity of a second TF to a nearby location. Such cooperative binding is important for activating gene expression from promoters and enhancers in both prokaryotic and eukaryotic cells. Existing methods to detect cooperative binding of a TF pair rely on analyzing the sequence that is bound. We propose a method that uses, instead, only ChIP-seq peak intensities and an expectation maximization (CPI-EM) algorithm. We validate our method using ChIP-seq data from cells where one of a pair of TFs under consideration has been genetically knocked out. Our algorithm relies on our observation that cooperative TF-TF binding is correlated with weak binding of one of the TFs, which we demonstrate in a variety of cell types, including *E. coli*, *S. cerevisiae*, *M. musculus*, as well as human cancer and stem cell lines. We show that this method performs significantly better than a predictor based only on the ChIP-seq peak distance of the TFs under consideration. This suggests that peak intensities contain information that can help detect the cooperative binding of a TF pair. The incorporation of peak intensities into existing sequence-based methods would allow them to detect new sequences capable of being cooperatively bound by TFs. The CPI-EM algorithm is available at <https://github.com/vishakad/cpi-em>.

1 Introduction

Transcription factors (TFs) regulate the transcription of a set of genes by binding specific regulatory regions of DNA. The magnitude of the change in transcription caused by a TF depends in part on its affinity to the DNA sequence bound. It is possible that a second TF binding a nearby sequence increases the first TF's binding affinity. In this case, the two TFs are said to cooperatively or combinatorially bind DNA [1]. The cooperative binding of transcription factors at enhancers and promoters is known to strongly increase gene expression [2, 3, 4, 5]. The presence of cooperativity has been used to explain the rapid rate of evolution of TF binding sites in multicellular organisms [6].

The role of cooperative binding in protein complex assembly has been extensively studied and computational methods have been proposed to detect such interactions within genomes [7, 8, 9]. In these studies, cooperativity typically involves protein oligomerization. However, two TFs can also cooperatively bind DNA without dimerizing prior to binding [10, 11]. Several theoretical methods have been proposed to detect cooperative binding between a pair of TFs in the genome [12, 13, 14, 15, 16, 17, 18, 19, 1]. These methods typically rely on locating frequently co-occurring binding sites of TF pairs across the genome, or within genomic sequences known to be bound by a TF pair. However, co-occurring binding site pairs do not always imply cooperative binding [1]. Conversely, many TF pairs can cooperatively bind DNA even if the spacing between their binding sites, or the sequence between them, is changed [20, 21]. Thus, a pair of TFs that cooperatively bind DNA at one genomic region may not bind cooperatively at a different genomic region.

^{*}To whom correspondence should be addressed. Email: vishakad@ncbs.res.in

33 Here, we propose a sequence-independent algorithm, based on ChIP-seq (chromatin immuno-precipitation and
34 sequencing) data, for detecting cooperatively bound sites that complement these sequence-based methods.

35 Genome-wide TF-DNA binding has been extensively studied using ChIP-seq [22]. ChIP-seq provides a list of
36 locations bound by a TF across a genome *in vivo*, which are referred to as *peaks*, along with peak *intensities*
37 whose values are proportional to the TF's affinity for the sequence bound at these locations [22]. Some ChIP-seq
38 experiments, designed to detect pair-wise TF interactions across the genome, have been carried out in *E. coli*, *S.*
39 *cerevisiae*, *M. musculus*, and human genomes [23, 1, 24, 25]. In these experiments, three sets of ChIP-seq are
40 performed to determine locations where a pair of TFs, A and B, are cooperatively bound. First, two ChIP-seq
41 experiments are performed to determine binding locations of A and B in cells. A third ChIP-seq is performed to find
42 binding locations of A after B is genetically knocked out. In this ChIP-seq, locations where A no longer binds DNA
43 or has a lower binding affinity towards DNA are considered to be instances of cooperative binding. We refer to such
44 a set of three experiments as A-B, and refer to A as the *primary* TF and B as the *partner* TF. Instead of knocking
45 out B, if a ChIP-seq is performed to find binding locations of B after A is knocked out, we can infer locations where
46 B is cooperatively bound by A. This dataset is labeled B-A, with B and A as primary and partner TFs, respectively.

47 We propose the ChIP-seq Peak Intensity - Expectation Maximisation (CPI-EM) algorithm as a computational
48 method to detect genomic locations cooperatively bound by a TF pair, based on their ChIP-seq peak intensities.
49 CPI-EM can do this without the need for ChIP-seq to be performed on one of the TFs after the other is knocked
50 out. At each location where ChIP-seq peaks of two TFs overlap each other, CPI-EM computes a probability that the
51 location is cooperatively bound by both TFs. The highlight of this algorithm is that it utilizes only peak intensities
52 to detect cooperative binding, and does not rely on binding site searches within ChIP-seq peak regions. CPI-EM
53 relies on the observation that a primary TF tended to be more weakly bound when it cooperatively bound DNA
54 with a partner TF, in comparison to regions where it did not cooperatively bind DNA. We observed this to be the
55 case in ChIP-seq datasets we analyzed from *E. coli*, *S. cerevisiae*, *M. musculus* and human genomes.

56 We compared the set of locations predicted by CPI-EM to be cooperatively bound, with the locations obtained
57 from the knockout-based ChIP-seq experiments. We also compared CPI-EM with an algorithm that detects coop-
58 erative binding based on the distance between ChIP-seq peaks. We found that peak distance by itself was not a
59 reliable predictor of cooperative binding. In contrast, we found that peak intensities are a more reliable criterion to
60 detect cooperative interactions in all the ChIP-seq datasets we analyzed.

61 2 Methods

62 2.1 ChIP-seq processing pipeline

63 A single ChIP-seq “peak call” consists of the genomic coordinates of the location being bound, along with a *peak*
64 *intensity*. We determined ChIP-seq peak locations of different transcription factors from multiple genomes, namely,
65 *E. coli* (GSE92255), *S. cerevisiae* [1], cells from primary *M. musculus* liver tissue [23], and three human cell lines
66 – the Caco-2 intestinal stem cell line [24], the T-47D breast cancer cell line and the ECC-1 endometrial cancer cell
67 line [25]. We used our own ChIP-seq pipeline to process raw sequence reads and call peaks from *M. musculus* and *S.*
68 *cerevisiae* data, and utilized pre-computed peak calls with the remaining datasets. This ensured that our validation
69 sets were not biased by procedures employed in our pipeline. See Supplementary Section 1 for details of our ChIP-seq
70 pipeline for processing these datasets.

71 2.2 Using ChIP-seq data from a genetic knockout to infer cooperative binding

72 From ChIP-seq profiles of a pair of TFs, X and Y, we classified genomic regions containing overlapping ChIP-seq
73 peaks of X and Y as cooperative or non-cooperative, based on the change in peak rank of X in response to a genetic
74 deletion of Y. The ranks are assigned such that the peak with rank 1 has the highest peak intensity. In our analysis,
75 we consider a genomic region to be doubly bound by X and Y if their peak regions overlap by at least a single base
76 pair. We used pybedtools v0.6.9 [26] to find these overlapping peak regions.

77 At each doubly bound genomic location, we classify X as being cooperatively bound by Y if (a) the peak rank
78 of X in the presence of Y is significantly higher (i.e., closer to rank 1) than the peak rank of X measured after the

79 deletion of Y, or (b) if X's peak is absent after the deletion of Y.

80 On the other hand, if the peak rank of X in the presence of Y is significantly lower (i.e., further from rank 1)
81 than the peak rank of X after the deletion of Y, or if it stays the same, we classify this as competitive or independent
82 binding, respectively. We refer to both these classes as non-cooperative binding. See Supplement Section 3 for
83 details on the statistical tests we performed to detect significant changes in peak ranks of X upon the knockout of
84 Y. These tests require ChIP-seq data from multiple replicates. In the ER α -FOXA1, CDX2-HNF4A, CRP-FIS, and
85 FIS-CRP datasets, peak calls from individual replicates were not available, therefore we used only peak losses to find
86 cooperatively bound locations in these datasets.

87 2.3 The ChIP-seq Peak Intensity - Expectation Maximisation (CPI-EM) algorithm

88 We describe the working of the CPI-EM algorithm in step-wise fashion below, where each of the steps is numbered
89 according to Figure 1. In Figure 1 and in the description below, we assume that cooperative binding between TFs
90 X and Y is being studied, where X is the primary TF and Y is the partner TF.

91 **Step 1:** From the ChIP-seq of X and Y, find all pairs of peaks where X and Y overlap by at least one base pair.
92 With these overlapping pairs, make a list of peak intensities $(x_1, y_1), (x_2, y_2) \dots (x_n, y_n)$, where x_i and y_i are the peak
93 intensities of the i -th peak of X and Y, respectively. This list of peak intensity pairs is the input data for the
94 CPI-EM algorithm.

95 **Step 2:** To this input data, fit a model of the joint probability $p(x, y)$ of observing the peak intensity x and y from
96 TFs X and Y, respectively, at a given location. Our model consists of a sum of two probability functions, which are
97 the probability of observing intensities x and y if they were (a) cooperatively bound, or (b) non-cooperatively bound.
98 We assume that both probability functions that are fitted have a Lognormal shape. This shape is characterized
99 by four parameters — a mean and a variance of the X and the Y axes (we also examine other shapes such as the
100 Gamma or Gaussian functions — see Supplementary Section 4). A final ninth parameter sets the relative weight of
101 the two probability functions, which determines the fraction of overlapping pairs that are cooperatively bound. We
102 find the best fit for these nine parameters using a procedure called expectation maximization (described in detail in
103 Supplementary Section 4.1).

104 We make two other assumptions in this step, each of which is discussed further in Supplementary Section 4.

- 105 • The peak intensities of X and Y at a location are statistically independent, irrespective of whether X and Y
106 are cooperatively or non-cooperatively bound. We found this to be a reasonable assumption after we measured
107 the mutual information between peak intensities of X and Y from cooperatively and non-cooperatively bound
108 locations. Mutual information is known to be a robust measure of statistical dependence [27].
- 109 • A primary TF that is cooperatively bound to DNA is, on average, bound weaker than a non-cooperatively
110 bound primary TF. We found this assumption to hold across all the datasets on which we ran CPI-EM (see
111 section “Peak intensities of cooperatively bound primary TFs are weaker than non-cooperatively bound primary
112 TFs” in Results, and Figure 2).

113 **Step 3:** Given the best-fit parameters, use Bayes' formula to calculate the probability for each overlapping pair
114 of ChIP-seq peaks to be a site of cooperative binding (see Supplementary Section 4).

115 **Step 4:** Choose a threshold probability α and label an overlapping pair as cooperatively bound if the probability
116 calculated in step 3 is greater than α , and as being non-cooperatively bound otherwise. Validate with a list of known
117 cooperative binding sites, e.g., derived from the ChIP-seq of X after Y is knocked out (as described in the previous
118 section).

119 2.4 Peak Distance Detector

120 For each peak intensity pair in the input data, the peak distance detector calculates the distance between the *summits*
121 of X and Y peak regions. The summit is a location within each peak region that has the highest number of sequence
122 reads that overlap it, and is typically the most likely site at which the TF is physically attached to DNA. The peak
123 distance detector declares doubly bound regions as cooperatively bound if the distance between peaks of X and Y
124 is lesser than a threshold distance d . We ran this detection algorithm on all the datasets on which CPI-EM was
125 employed to detect cooperative binding. Our goal in using this algorithm was to determine whether the distance
126 between peaks is a reliable criterion to discriminate between cooperative and non-cooperative binding.

3 Results

3.1 Peak intensities of cooperatively bound primary TFs are weaker than non-cooperatively bound primary TFs

We inferred cooperative and non-cooperative binding from ChIP-seq datasets of FIS-CRP and CRP-FIS pairs in *E. coli* in early-exponential and mid-exponential growth phases (accession number GSE92255), GCN4-RTG3 and RTG3-GCN4 in *S. cerevisiae* [1], FOXA1-HNF4A, FOXA1-CEBPA and HNF4A-CEBPA in the mouse (*M. musculus*) liver [23], CDX2-HNF4A in differentiated human intestinal stem cell lines (Caco-2) [24], and ER α -FOXA1 in the T-47D breast cancer cell line [25]. A summary of the data is shown in Supplementary Table 1.

Figure 2A-C summarize trends in cooperative and non-cooperative TF-DNA binding seen in these datasets. Cooperatively and non-cooperatively bound locations were determined using ChIP-seq data from genetic knockouts as discussed in Methods. Cooperatively bound primary TF peak intensities were significantly lower than those of non-cooperatively bound primary TF peaks across each of the TF-TF pairs (Wilcoxon rank-sum test, $p \ll 0.001$). In contrast, there was no consistent trend in the intensities of the partner TF in each of these pairs. This meant that a primary TF could be cooperatively bound to DNA irrespective of the peak intensity of the partner TF. In Figure 2B, estimates of the marginal distributions of cooperative and non-cooperative FOXA1 and HNF4A peaks are shown. These, and similar distributions for the other TF pairs, tended to be better approximated by a Lognormal distribution, which was evident from the higher log-likelihood value associated with a Lognormal fit, compared to a Gaussian or Gamma distribution (Supplementary Table 2).

Since the primary TF intensity distributions from cooperatively bound regions significantly differed from those of non-cooperatively bound regions, it should be possible to accurately label a pair of overlapping peaks as cooperative or non-cooperative, based solely on their peak intensities. For instance, in the FOXA1-HNF4A dataset, a FOXA1 peak that has an intensity value of 5 is ≈ 3.4 times more likely to be cooperatively bound with HNF4A than to be non-cooperatively bound with it. In clear-cut cases such as these, knowledge of the underlying sequence that is bound is not necessary to detect a cooperative interaction.

3.2 CPI-EM applied to ChIP-seq datasets from *M. musculus*, *S. cerevisiae* and *E. coli*

The ChIP-seq Peak Intensity - Expectation Maximisation (CPI-EM) algorithm works as illustrated in Figure 1 (with a detailed explanation in the Methods).

Figure 3 shows the result of the CPI-EM algorithm when used to predict genomic regions that are cooperatively bound by FOXA1-HNF4A, RTG3-GCN4 and FIS-CRP in *M. musculus*, *S. cerevisiae* and early-exponential phase cultures of *E. coli*, respectively. The top row shows histograms of the cooperative binding probabilities ($p_1^{coop}, p_2^{coop}, \dots, p_N^{coop}$), which are computed by CPI-EM, for all peak intensity pairs from each of the three datasets. The height of each bar is the fraction of peak intensity pairs in each probability bin that are actually cooperatively bound (termed true positives, which are calculated based on knockout data as explained in Methods). True positives are distributed differently between the bins across different datasets. Over 50% of cooperatively bound RTG3-GCN4 pairs are assigned a value of $p_{coop} > 0.8$ by CPI-EM, with over 90% of cooperative bound pairs having a value of $p_{coop} > 0.5$. In contrast, only about 30% of cooperative FOXA1-HNF4A pairs have a $p_{coop} > 0.8$.

The distribution of cooperative pairs into each of these bins determines the number of errors made when all peak pairs with $p_{coop} > \alpha$ are declared as cooperatively bound. The false positive rate (FPR) of the CPI-EM algorithm is the fraction of non-cooperatively bound regions erroneously declared as cooperatively bound, while the true positive rate (TPR) is the fraction of cooperatively bound regions that are detected. Both these quantities are functions of α , and are estimated as

$$FPR(\alpha) = \frac{N_{FP}(\alpha)}{N_{nc}}, \quad TPR(\alpha) = \frac{N_{TP}(\alpha)}{N_c},$$

where $N_{FP}(\alpha)$ is the number of non-cooperatively bound regions mistakenly declared as cooperatively bound at a threshold α , while $N_{TP}(\alpha)$ is the number of cooperatively bound regions correctly declared as cooperatively bound with the threshold α . N_c and N_{nc} represent the total number of cooperatively bound and non-cooperatively bound regions, respectively. The receiver operating characteristic (ROC) curves at the bottom row of Figure 3 shows the trade-off between *false positive rates* and *true positive rates* of CPI-EM at different values of α . A larger value of α results in fewer false positives in the final prediction set but also results in fewer true positives being detected. For instance, in the CRP-FIS dataset, $\alpha = 0.81$ allows nearly 50% of all cooperative interactions to be detected. If α is lowered to 0.375, more than 75% of cooperative peak pairs can be detected, but there will be more false positives in this prediction set since the FPR at this value of α is three times higher than that at $\alpha = 0.81$. A way of quantifying

173 the detection performance of an algorithm like CPI-EM is to calculate the area under the ROC curve (auROC).
174 This is a measure of the average true positive rate of the CPI-EM algorithm, with a higher value representing better
175 detection. Consequently, auROC also provides a way of comparing two different detection algorithms.

176 3.3 Performance of different variants of the CPI-EM algorithm

177 When CPI-EM was run to compute cooperative binding probabilities in Figure 3, Lognormal shapes were fitted to
178 the joint probability function of observing cooperative and non-cooperatively bound peak intensity pairs (step 2 of
179 the CPI-EM algorithm). In Figure 4, we compare the detection performance of this version of the CPI-EM algorithm
180 with two other variants that fit Gamma and Gaussian shapes instead of a Lognormal shape. Figure 4 shows the
181 auROC of these three variants of the CPI-EM algorithm after they were run on all the datasets shown in Figure
182 2. We also compared the auROC of these CPI-EM variants to the auROC of a “naive” peak distance detector
183 and a detector based purely on chance. The peak distance algorithm computes the distances between the peaks of
184 overlapping ChIP-seq peaks and declares those overlapping peak pairs whose peaks are within a threshold distance
185 d to be cooperatively bound. The chance detector is based on using tosses from a biased coin, whose probability of
186 showing heads is α , to detect cooperative interactions. The area under the ROC of this detector will be 0.5 for any
187 dataset (see Supplementary Section 5). An auROC of 0.5 thus represents the minimum level of detection performance
188 that an algorithm should obtain to be considered a useful detector in practice. The complete ROC curves of each of
189 the CPI-EM and peak distance algorithms for the datasets in Figure 2 are shown in Supplementary Figure 1.

190 In Figure 4, it can be seen that the Gamma and Lognormal CPI-EM variants have an auROC of at least 0.5, and
191 thus can consistently detect cooperative interactions across all datasets. The Lognormal CPI-EM variant fares well
192 on all datasets, except for the mid-exponential phase CRP-FIS dataset, where its performance is at the level of a
193 chance detector. The Gaussian CPI-EM variant performs poorly on the early-exponential phase FIS-CRP and CRP-
194 FIS datasets and has an auROC less than that of a chance detector in the former. This indicates that the Gaussian
195 CPI-EM variant is not as reliable as the Gamma and Lognormal variants in detecting cooperative interactions.

196 There is considerable variation in the auROC of the peak distance based algorithm: less than 0.5 in early-
197 exponential phase CRP-FIS and RTG3-GCN4 datasets, but higher than 0.5 in the remaining datasets. The fact
198 that this algorithm can perform worse than a chance detector shows that peak distance, by itself, is an unreliable
199 criterion for detecting cooperative binding. This is in contrast to the reliable (auROC > 0.7 across most datasets)
200 performance obtained with the Lognormal and Gamma CPI-EM algorithms.

201 3.4 Application of CPI-EM to detect cell-type specific binding of transcription factors

202 We now demonstrate an application of CPI-EM in detecting the cell-type specific binding of a TF in two ChIP-
203 seq data sets from cell lines. Many studies of cell-type specific binding are targeted at understanding cellular
204 reprogramming and stem cell differentiation that gives rise to various organs in animal development [10]. One of the
205 reasons the same TF can bind different genomic regions in two cell types is because it cooperatively binds DNA with
206 the lineage-determining transcription factor of that cell type [10]. Thus, the lineage-determining transcription factor
207 effectively functions as a partner TF, whose concentration is low in one cell type but high in a second one.

208 We ran CPI-EM on ChIP-seq datasets of ER α -FOXA1 from T-47D breast cancer and ECC-1 endometrial cancer
209 cell lines [25], and CDX2-HNF4A in proliferating and differentiated Caco-2 human intestinal stem cell lines [24].
210 In the ER α -FOXA1 dataset, we compared ER α binding between T-47D and ECC-1 cell lines. T-47D cells express
211 FOXA1 at a ≈ 50 fold higher concentration than ECC-1 cell lines, with this difference in concentration correlated
212 with differences in ER α occupancy [25]. Given such a large difference in FOXA1 concentration between these two
213 cell types, we treated the ChIP-seq of ER α in ECC-1 cell lines as being equivalent to a knockout of FOXA1 from
214 T-47D cells. Similarly, in the CDX2-HNF4A dataset, differentiated Caco-2 cells express HNF4A at a much higher
215 concentration than proliferating Caco-2 cells [24]. Thus, a comparison of CDX2 binding between differentiated
216 and proliferating Caco-2 cells is akin to analyzing changes in CDX2 binding after HNF4A is knocked out from
217 differentiated Caco-2 cells. Further, CDX2 has been shown to cooperatively bind DNA with HNF4A through an
218 independent biochemical assay [24]. However, the loss of binding in one cell type compared to the other is not
219 solely due to cooperative binding with FOXA1 or HNF4A—differences in nucleosomal occupancy and modifications
220 between cell types are known to influence cell-type specific binding in both datasets [25, 24]. Nevertheless, we wanted
221 to see if CPI-EM could still detect cell-type specific binding in both these datasets.

222 The box plots in Figure 5A compare the distribution of intensities of ER α peaks present only in T-47D cells
223 with that of ER α peaks present in both cell types. Although other factors determine cell-type specific binding of
224 ER α , we found that ER α peaks present only in T-47D cells were of lower intensity than peaks present in both cell
225 types (Wilcoxon rank-sum test, $p \ll 0.001$). The same trend was seen in CDX2-HNF4A, where regions occupied

226 by CDX2 only in differentiated Caco-2 cells were more weakly bound than regions occupied by CDX2 in both cell
227 types (Wilcoxon rank-sum test, $p \ll 0.001$). In contrast, the trends in the intensities of the partner TFs FOXA1 and
228 HNF4A in both these datasets are different. FOXA1 peaks in ER α cell-type specific bound regions are actually more
229 strongly bound than FOXA1 peaks in shared ER α bound regions, while HNF4A peaks in CDX2 cell-type specific
230 bound regions are more weakly bound than in shared CDX2 bound regions.

231 Since these patterns in peak intensities of cell-type specific binding were similar to those of cooperative binding
232 seen in Figure 2, we ran all three variants of the CPI-EM algorithm on ER α -FOXA1 and CDX2-HNF4A data sets
233 (Figure 5B). The differences in performance between the Gamma and Lognormal CPI-EM variants are marginal but
234 they are both better than the Gaussian variant at detecting cell-type specific binding in these two datasets. The
235 peak distance detector could not be tested on ER α -FOXA1 data since peak locations were not available in the peak
236 calls. In the CDX2-HNF4A dataset, however, the peak distance algorithm has an auROC only marginally higher
237 than 0.5. Once again, the peak distance criterion is poor at detecting cooperative binding in this dataset. This is in
238 contrast to the CPI-EM algorithm, which has an auROC greater than 0.5 in both datasets, with the Gamma and
239 Lognormal CPI-EM variants giving an auROC of 0.71 in the CDX2-HNF4A dataset.

240 4 Discussion

241 Cooperative binding is known to play a role in transcription factor binding site evolution and enhancer detection
242 [28]. Cooperativity is also known to influence cis-regulatory variation between individuals of a species [29], which
243 could potentially capture disease-causing mutations that are known to occur in regulatory regions of the genome
244 [30]. CPI-EM is suited to study these phenomena since it can detect instances of cooperative binding between a pair
245 of transcription factors that may occur anywhere in the genome. While sequence-based approaches to cooperative
246 binding detection have been proposed [12, 13, 14, 15, 16, 17, 18, 19, 1], none use ChIP-seq peak intensities as a
247 criterion to detect cooperativity. Our goal was to demonstrate that peak intensities, by themselves, contain valuable
248 information to detect cooperative binding. Our results suggest that methods for detecting cooperative binding based
249 on ChIP-seq peak intensities can usefully complement sequence-based detection algorithms.

250 4.1 Assumptions in the CPI-EM algorithm

251 The assumption that cooperatively bound primary TFs are more weakly bound, on average, than non-cooperatively
252 bound primary TFs is the key assumption in the CPI-EM algorithm. This assumption was true across TF pairs we
253 analyzed in *E. coli*, *S. cerevisiae*, *M. musculus*, and human genomes (which included cancer and stem cell lines). A
254 consequence of this assumption is that CPI-EM is unlikely to detect regions where the primary TF is cooperatively
255 bound to DNA, but with a high peak intensity. Further, a fraction of low intensity peaks are known to correspond
256 to *indirect* or piggy-back binding of a TF to DNA, where the TF does not directly bind DNA but instead binds
257 a second protein that in turn binds DNA [25, 35]. However, sequence-based methods to detect such interactions
258 [36, 37] can be used to filter out such ChIP-seq peaks before they are input to the CPI-EM algorithm.

259 Our observation that cooperatively bound TFs were more weakly bound than non-cooperatively bound TFs is
260 likely a signature of short-range pair-wise interactions. For instance, GCN4-RTG3 and CDX2-HNF4A interactions
261 were discovered in the datasets upon which we ran CPI-EM, and these interactions have been independently verified
262 [1, 24]. A similar pattern of weakly but cooperatively bound TFs is seen in animal development. The binding of
263 Ultrabithorax (Ubx) and Extradenticle at the *shavenbaby* enhancer in *Drosophila melanogaster* embryos [31] occurs
264 in closely spaced low-affinity binding sites to help coordinate tissue patterning. Mutations that increased Ubx binding
265 affinity led to the expression of proteins outside their naturally occurring tissue boundaries [31]. Similarly, low-affinity
266 binding sites that cooperatively bind Cubitus interruptus at the *dpp* enhancer (which plays a crucial role in wing
267 patterning in *Drosophila melanogaster*) are evolutionary conserved across twelve *Drosophila* species [32].

268 4.2 Challenges to cooperativity detection using ChIP-seq peak intensities

269 There are two principal challenges to detecting cooperative interactions using ChIP-seq peak intensities — its low
270 spatial resolution, and the use of PCR amplification. ChIP-seq cannot resolve binding events that occur within 100
271 base pairs of each other [33], while DNA-mediated cooperative binding often occurs between two TFs bound within
272 25 base pairs of each other [11, 20]. Thus, a single ChIP-seq peak intensity can represent the average of multiple
273 cooperative and non-cooperative binding events. This low resolution may also explain why peak distance was not a
274 reliable criterion to detect cooperative binding. Protocols such as ChIP-exo [33] and ChIP-nexus [34] can resolve two
275 binding events that are a single base pair apart [35]. These methods likely provide more accurate measurements of

276 distances between binding events, which means that ChIP-exo or ChIP-nexus peak distances may supplement peak
277 intensities in detecting cooperative interactions.

278 Peak intensities are also affected by PCR amplification, which is a necessary step in ChIP-seq protocols. While
279 the use of more PCR amplification cycles helps detect weaker binding events, the variance in the number of fragments
280 obtained at the end of the PCR process increases with the number of cycles employed [38, 39]. If peak intensities
281 can instead be calculated based on the number of *un-amplified* DNA fragments, they would be less noisy measures
282 of binding affinity. This is possible with protocols such as ChIP-nexus [34], that use molecular bar-coding techniques
283 in DNA library preparation [40].

284 The additional variance introduced by PCR amplification might also explain the low mutual information values
285 we measured between peak intensities of cooperatively bound TFs (see Supplementary Table 2). Thus, protocols
286 such as ChIP-nexus and ChIP-exo might be sensitive enough to detect the difference in MI between cooperatively
287 and non-cooperatively bound TFs [41]. In such a case, our method can be modified to no longer be dependent on the
288 assumption of cooperatively bound primary TFs being more weakly bound than non-cooperatively bound primary
289 TFs. In this modified algorithm, a tenth parameter in the joint probability model fit to peak intensity data (in step
290 2 of the CPI-EM algorithm) will take into account this mutual information resulting from cooperative binding. The
291 precise form of such a modified joint probability model is not obvious, but it would increase the probability that a
292 high MI peak intensity pair would be labeled as cooperative, despite having a strongly bound primary TF.

293 Ultimately, our method is a way of detecting cooperatively bound locations without making any direct assumptions
294 about the genomic sequence of that location. Therefore, it provides a useful way of finding binding sequence patterns
295 that allow for cooperative binding to occur *in vivo* but lie outside the range of existing sequence-based algorithms.

296 5 Funding

297 Support from the Simons Foundation (to S.K. and V.D.); PRISM 12th plan project at Institute of Mathematical
298 Sciences (to R.S.);

299 6 Acknowledgements

300 We thank Aswin Sai Narain Seshasayee, Parul Singh, Vijay Kumar, Deepa Agashe, and Leelavati Narlikar for
301 discussions.

302 *Author Contributions* : V.D. conceived the study, and designed and implemented the CPI-EM algorithm. V.D.,
303 R.S., and S.K. analyzed and interpreted the results, and wrote the manuscript.

304 References

- 305 [1] Spivak, A. T. and Stormo, G. D. . Combinatorial cis-regulation in *saccharomyces* species. *G3: Genes—*
306 *Genomes— Genetics*, 6(3):653–667, 2016.
- 307 [2] Patwardhan, R. P. , Hiatt, J. B. , Witten, D. M. , Kim, M. J. , Smith, R. P. , May, D. , Lee, C. , Andrie, J. M.
308 , Lee, S.-I. , Cooper, G. M. , et al. Massively parallel functional dissection of mammalian enhancers *in vivo*.
309 *Nature biotechnology*, 30(3):265–270, 2012.
- 310 [3] Smith, R. P. , Taher, L. , Patwardhan, R. P. , Kim, M. J. , Inoue, F. , Shendure, J. , Ovcharenko, I. , and
311 Ahituv, N. . Massively parallel decoding of mammalian regulatory sequences supports a flexible organizational
312 model. *Nature Genetics*, 45(9):1021–1028, 2013.
- 313 [4] Sharon, E. , Kalma, Y. , Sharp, A. , Raveh-Sadka, T. , Levo, M. , Zeevi, D. , Keren, L. , Yakhini, Z. ,
314 Weinberger, A. , and Segal, E. . Inferring gene regulatory logic from high-throughput measurements of thousands
315 of systematically designed promoters. *Nature Biotechnology*, 30(6):521–530, 2012.
- 316 [5] Gutierrez, P. S. , Monteoliva, D. , and Diambra, L. . Cooperative binding of transcription factors promotes
317 bimodal gene expression response. *PLoS One*, 7(9):e44812, 2012.
- 318 [6] Tuğrul, M. , Paixão, T. , Barton, N. H. , and Tkačik, G. . Dynamics of transcription factor binding site evolution.
319 *PLoS Genet*, 11(11):e1005639, 2015.

- 320 [7] Jansen, R. , Yu, H. , Greenbaum, D. , Kluger, Y. , Krogan, N. J. , Chung, S. , Emili, A. , Snyder, M. ,
321 Greenblatt, J. F. , and Gerstein, M. . A bayesian networks approach for predicting protein-protein interactions
322 from genomic data. *science*, 302(5644):449–453, 2003.
- 323 [8] Krogan, N. J. , Cagney, G. , Yu, H. , Zhong, G. , Guo, X. , Ignatchenko, A. , Li, J. , Pu, S. , Datta, N. ,
324 Tikuisis, A. P. , et al. Global landscape of protein complexes in the yeast *saccharomyces cerevisiae*. *Nature*,
325 440(7084):637–643, 2006.
- 326 [9] Hardison, R. . Hemoglobins from bacteria to man: evolution of different patterns of gene expression. *Journal*
327 *of Experimental Biology*, 201(8):1099–1117, 1998.
- 328 [10] Reiter, F. , Wienerroither, S. , and Stark, A. . Combinatorial function of transcription factors and cofactors.
329 *Current Opinion in Genetics & Development*, 43:73–81, 2017.
- 330 [11] Kim, S. , Broströmer, E. , Xing, D. , Jin, J. , Chong, S. , Ge, H. , Wang, S. , Gu, C. , Yang, L. , Gao, Y. Q. ,
331 et al. Probing allostery through DNA. *Science*, 339(6121):816–819, 2013.
- 332 [12] GuhaThakurta, D. and Stormo, G. D. . Identifying target sites for cooperatively binding factors. *Bioinformatics*,
333 17(7):608–621, 2001.
- 334 [13] Whittington, T. , Frith, M. C. , Johnson, J. , and Bailey, T. L. . Inferring transcription factor complexes from
335 chip-seq data. *Nucleic Acids Research*, 39(15):e98–e98, 2011.
- 336 [14] Kazemian, M. , Pham, H. , Wolfe, S. A. , Brodsky, M. H. , and Sinha, S. . Widespread evidence of cooperative
337 dna binding by transcription factors in drosophila development. *Nucleic Acids Research*, 41(17):8237–8252,
338 2013.
- 339 [15] Das, D. , Banerjee, N. , and Zhang, M. Q. . Interacting models of cooperative gene regulation. *Proceedings of*
340 *the National Academy of Sciences of the United States of America*, 101(46):16234–16239, 2004.
- 341 [16] He, X. , Chen, C.-C. , Hong, F. , Fang, F. , Sinha, S. , Ng, H.-H. , and Zhong, S. . A biophysical model for
342 analysis of transcription factor interaction and binding site arrangement from genome-wide binding data. *PloS*
343 *One*, 4(12):e8155, 2009.
- 344 [17] Girgis, H. Z. and Ovcharenko, I. . Predicting tissue specific cis-regulatory modules in the human genome using
345 pairs of co-occurring motifs. *BMC Bioinformatics*, 13(1):1, 2012.
- 346 [18] Nandi, S. , Blais, A. , and Ioshikhes, I. . Identification of cis-regulatory modules in promoters of human genes
347 exploiting mutual positioning of transcription factors. *Nucleic Acids Research*, page gkt578, 2013.
- 348 [19] Jiang, P. and Singh, M. . Ccat: combinatorial code analysis tool for transcriptional regulation. *Nucleic Acids*
349 *Research*, 42(5):2833–2847, 2014.
- 350 [20] Jolma, A. , Yin, Y. , Nitta, K. R. , Dave, K. , Popov, A. , Taipale, M. , Enge, M. , Kivioja, T. , Morgunova, E. ,
351 and Taipale, J. . DNA-dependent formation of transcription factor pairs alters their binding specificity. *Nature*,
352 527(7578):384–388, 2015.
- 353 [21] Narasimhan, K. , Pillay, S. , Huang, Y.-H. , Jayabal, S. , Udayasuryan, B. , Veerapandian, V. , Kolatkar, P. ,
354 Cojocar, V. , Pervushin, K. , and Jauch, R. . DNA-mediated cooperativity facilitates the co-selection of cryptic
355 enhancer sequences by SOX2 and PAX6 transcription factors. *Nucleic Acids Research*, page gku1390, 2015.
- 356 [22] Johnson, D. S. , Mortazavi, A. , Myers, R. M. , and Wold, B. . Genome-wide mapping of in vivo protein-dna
357 interactions. *Science*, 316(5830):1497–1502, 2007.
- 358 [23] Stefflova, K. , Thybert, D. , Wilson, M. D. , Streeter, I. , Aleksic, J. , Karagianni, P. , Brazma, A. , Adams,
359 D. J. , Talianidis, I. , Marioni, J. C. , et al. Cooperativity and rapid evolution of cobound transcription factors
360 in closely related mammals. *Cell*, 154(3):530–540, 2013.
- 361 [24] Verzi, M. P. , Shin, H. , He, H. H. , Sulahian, R. , Meyer, C. A. , Montgomery, R. K. , Fleet, J. C. , Brown, M.
362 , Liu, X. S. , and Shivdasani, R. A. . Differentiation-specific histone modifications reveal dynamic chromatin
363 interactions and partners for the intestinal transcription factor *cdx2*. *Developmental Cell*, 19(5):713–726, 2010.

- 364 [25] Gertz, J. , Savic, D. , Varley, K. E. , Partridge, E. C. , Safi, A. , Jain, P. , Cooper, G. M. , Reddy, T. E. ,
365 Crawford, G. E. , and Myers, R. M. . Distinct properties of cell-type-specific and shared transcription factor
366 binding sites. *Molecular cell*, 52(1):25–36, 2013.
- 367 [26] Dale, R. K. , Pedersen, B. S. , and Quinlan, A. R. . Pybedtools: a flexible python library for manipulating
368 genomic datasets and annotations. *Bioinformatics*, 27(24):3423–3424, 2011.
- 369 [27] Kinney, J. B. and Atwal, G. S. . Equitability, mutual information, and the maximal information coefficient.
370 *Proceedings of the National Academy of Sciences*, 111(9):3354–3359, 2014.
- 371 [28] Villar, D. , Flicek, P. , and Odom, D. T. . Evolution of transcription factor binding in metazoans-mechanisms
372 and functional implications. *Nature Reviews Genetics*, 15(4):221–233, 2014.
- 373 [29] Heinz, S. , Romanoski, C. , Benner, C. , Allison, K. , Kaikkonen, M. , Orozco, L. , and Glass, C. . Effect of
374 natural genetic variation on enhancer selection and function. *Nature*, 503(7477):487–492, 2013.
- 375 [30] Knight, J. C. . Regulatory polymorphisms underlying complex disease traits. *Journal of Molecular Medicine*,
376 83(2):97–109, 2005.
- 377 [31] Crocker, J. , Abe, N. , Rinaldi, L. , McGregor, A. P. , Frankel, N. , Wang, S. , Alsawadi, A. , Valenti, P. , Plaza,
378 S. , Payre, F. , et al. Low affinity binding site clusters confer hox specificity and regulatory robustness. *Cell*,
379 160(1):191–203, 2015.
- 380 [32] Ramos, A. I. and Barolo, S. . Low-affinity transcription factor binding sites shape morphogen responses and
381 enhancer evolution. *Phil. Trans. R. Soc. B*, 368(1632):20130018, 2013.
- 382 [33] Rhee, H. S. and Pugh, B. F. . Comprehensive genome-wide protein-dna interactions detected at single-nucleotide
383 resolution. *Cell*, 147(6):1408–1419, 2011.
- 384 [34] He, Q. , Johnston, J. , and Zeitlinger, J. . Chip-nexus enables improved detection of in vivo transcription factor
385 binding footprints. *Nature Biotechnology*, 33(4):395–401, 2015.
- 386 [35] Starick, S. R. , Ibn-Salem, J. , Jurk, M. , Hernandez, C. , Love, M. I. , Chung, H.-R. , Vingron, M. , Thomas-
387 Chollier, M. , and Meijnsing, S. H. . Chip-exo signal associated with dna-binding motifs provides insight into
388 the genomic binding of the glucocorticoid receptor and cooperating transcription factors. *Genome Research*,
389 25(6):825–835, 2015.
- 390 [36] Gordán, R. , Hartemink, A. J. , and Bulyk, M. L. . Distinguishing direct versus indirect transcription factor-
391 DNA interactions. *Genome research*, 19(11):2090–2100, 2009.
- 392 [37] Bailey, T. L. and Machanick, P. . Inferring direct DNA binding from ChIP-seq. *Nucleic acids research*, page
393 gks433, 2012.
- 394 [38] Kebschull, J. M. and Zador, A. M. . Sources of PCR-induced distortions in high-throughput sequencing data
395 sets. *Nucleic Acids Research*, 43(21):e143–e143, 2015.
- 396 [39] Best, K. , Oakes, T. , Heather, J. M. , Shawe-Taylor, J. , and Chain, B. . Computational analysis of stochastic
397 heterogeneity in PCR amplification efficiency revealed by single molecule barcoding. *Scientific Reports*, 5, 2015.
- 398 [40] Kivioja, T. , Vähärautio, A. , Karlsson, K. , Bonke, M. , Enge, M. , Linnarsson, S. , and Taipale, J. . Counting
399 absolute numbers of molecules using unique molecular identifiers. *Nature methods*, 9(1):72–74, 2012.
- 400 [41] Kinney, J. B. , Tkačik, G. , and Callan, C. G. . Precise physical models of protein-dna interaction from
401 high-throughput data. *Proceedings of the National Academy of Sciences*, 104(2):501–506, 2007.

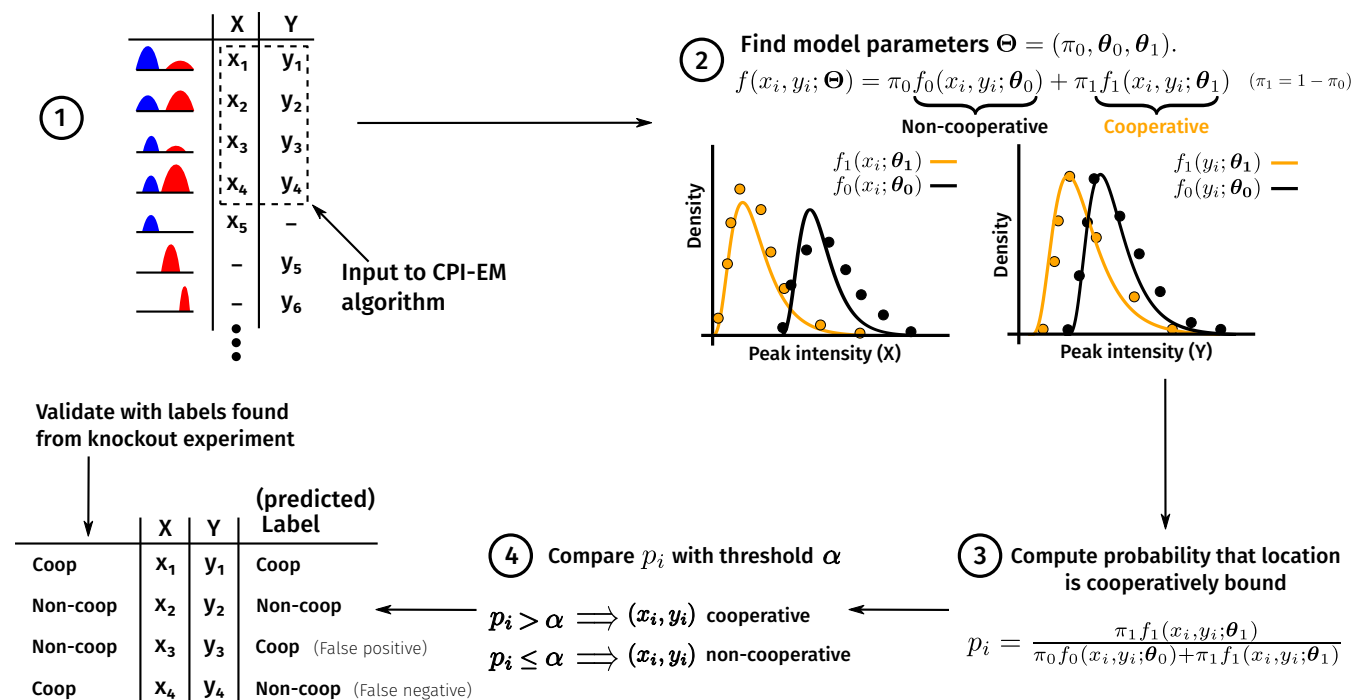


Figure 1: **A** schematic of the use of the CPI-EM algorithm and ChIP-seq from knockout data to separately identify cooperatively bound transcription factor pairs. ChIP-seq experiments carried out on two TFs, X and Y, yield a list of locations that are bound by both TFs, along with peak intensities at each location. From this data, there are two ways in which we find genomic locations that are cooperatively bound by X and Y.

(A) A method for inferring these locations from a ChIP-seq of X carried out after Y is genetically deleted. Locations where a peak of X either disappears altogether, or is reduced in intensity after knocking out Y are labelled as cooperatively bound. In contrast, locations where a peak of X either remains unchanged or increases in intensity are labelled as non-cooperatively bound (see section “Using ChIP-seq data from a genetic knockout to infer cooperative binding” in Methods).

(B) Steps in predicting cooperatively bound locations are shown, where the numbers correspond to those in the section “The ChIP-seq Peak Intensity - Expectation Maximisation (CPI-EM) algorithm” in Methods. (1) The input to CPI-EM consists of a list of genomic locations where a peak of X overlaps a peak of Y by at least a single base pair. (2) Each of these overlapping intensity pairs is fit to a model that consists of a sum of two probability functions. These functions specify the probabilities of observing a particular peak intensity pair given that it comes from a cooperatively or non-cooperatively bound region. These probabilities are computed by fitting the model to the input data using the expectation-maximization algorithm (see Supplementary Section 4.1). (3) Bayes’ formula is applied to the probabilities computed in step (2) to find the probability of each peak intensity pair being cooperatively bound. (4) Each cooperative binding probability computed in step (3) that is greater than a threshold α is declared as cooperatively bound. We compare this list of predicted locations with the list of cooperatively bound locations inferred from knockout data in order to compute the number of correct and incorrect inferences made by CPI-EM.

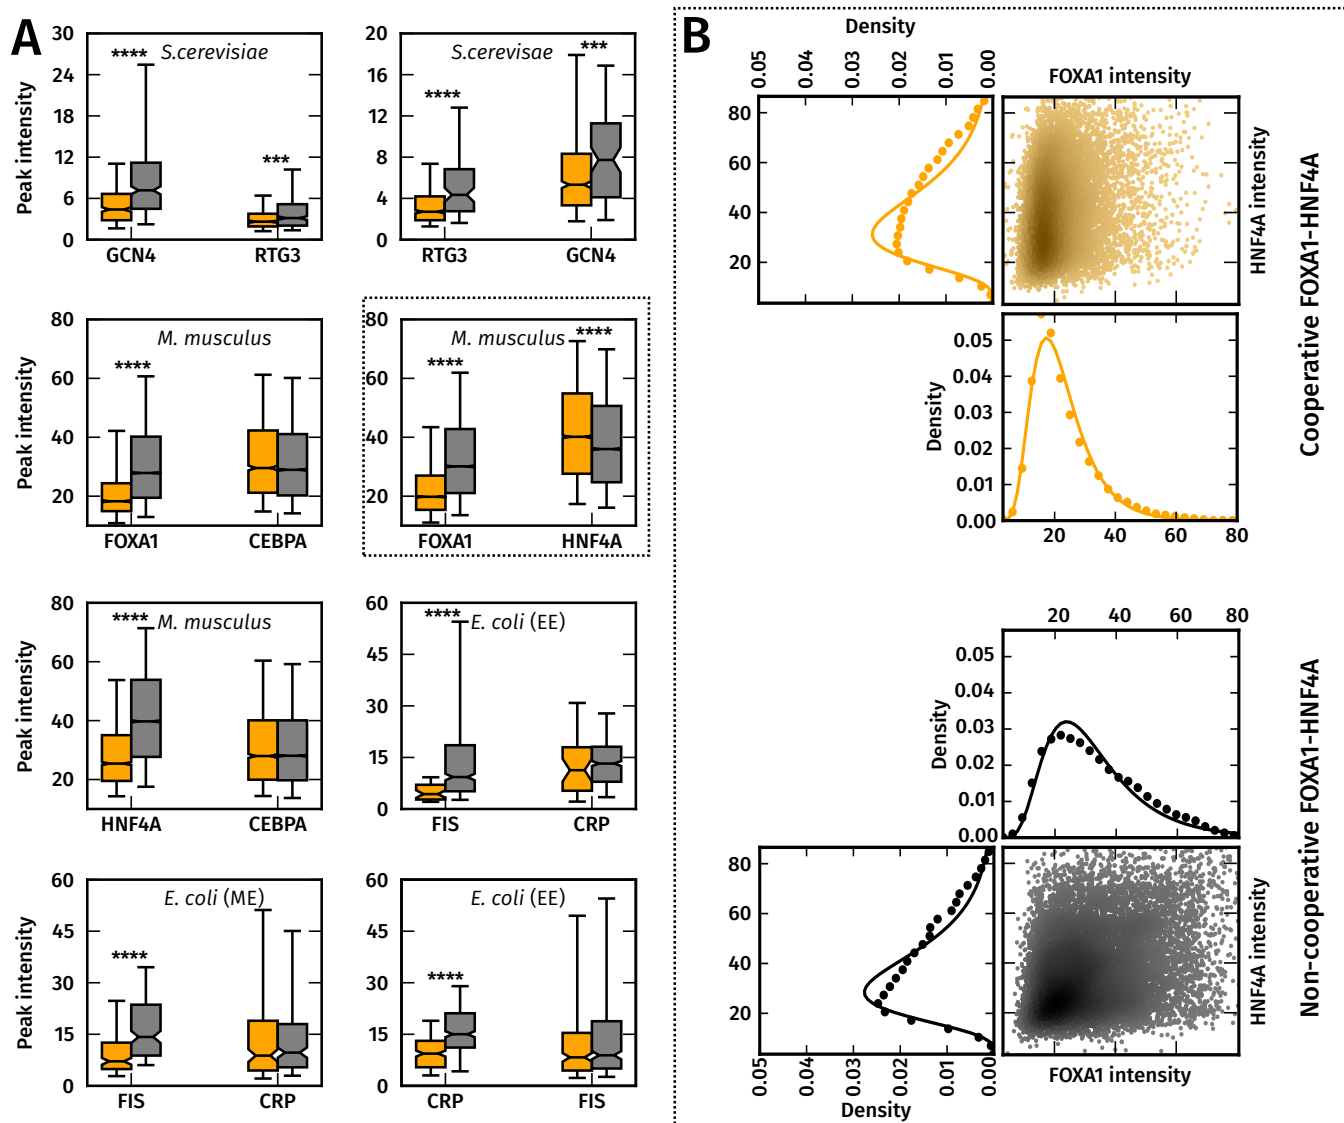


Figure 2: (A) **Cooperatively bound primary TFs are significantly more weakly bound than non-cooperatively bound primary TFs.** Box-plots of peak intensity distributions of cooperatively (orange) and non-cooperatively (gray) bound TF pairs, with primary TFs on the left and partner TFs on the right. ****, *** and ** indicate p-values of $< 10^{-4}$, 10^{-3} and 10^{-2} from a Wilcoxon rank sum test. The whiskers of the box plot are the 5 - th and 95 - th percentiles of the distributions shown.

(B) **ChIP-seq peak intensity distributions can be approximated by a Lognormal distribution.** Marginal peak intensity distributions of FOXA1 and HNF4A peaks (in filled black and orange circles), with fitted Lognormal distributions (solid black and orange lines), along side a scatter plot of (FOXA1,HNF4A) peak intensity pairs from cooperatively and non-cooperatively bound regions. The scatter points are colored according to the density of points in that region, with darker shades indicating a higher density. The density of points in the scatter were computed using the Gaussian kernel density estimation procedure in the Python Scipy library.




Figure 3: **CPI-EM applied to ChIP-seq datasets from *M. musculus* (FOXA1-HNF4A), *S. cerevisiae* (RTG3-GCN4) and early-exponential phase cultures of *E. coli* (CRP-FIS).** For each dataset, CPI-EM computes a list of cooperative binding probabilities at all the locations bound by the TF pair under consideration. **Top row:** The fraction of cooperatively bound pairs, as determined from knockout data, that fall into each cooperative binding probability bin. The bins are equally spaced with a width of 0.1 and the heights of the bars within each histogram add up to 1. **Bottom row: Receiver operating characteristic (ROC) curves that evaluate the performance of CPI-EM in detecting cooperatively bound pairs.** The curve is generated by calculating, for each value of α between 0 and 1, the true and false positive rate of the algorithm. The true positive rate ($TPR(\alpha)$) is the ratio of the number of cooperatively bound regions detected to the total number of cooperatively bound regions at that value of α . The false positive rate ($FPR(\alpha)$) is the ratio of the number of non-cooperatively bound regions mistakenly detected as cooperatively bound to the total number of non-cooperatively bound regions at that value of α . Small values of α give a higher TPR, but at the cost of a higher FPR. The area under the ROC (auROC) is a measure of detection performance, whose value cannot exceed 1, which corresponds to a perfect detector. Given the auROC of two different algorithms, the one with a higher auROC is better, on average, at detecting cooperative binding.




Figure 4: **The CPI-EM variant that fits lognormal distributions to peak intensity pairs consistently performs well across all datasets** The area under the ROC curve (auROC) of the CPI-EM algorithm applied to each of the datasets shown in Figure 2. CPI-EM variants that fit Lognormal, Gamma and Gaussian distributions are represented in orange, black and gray, respectively. The auROC of the peak distance based detector is shown in blue. See Supplementary Section 5 for the calculation of the ROC curve for both the CPI-EM and peak distance algorithms.

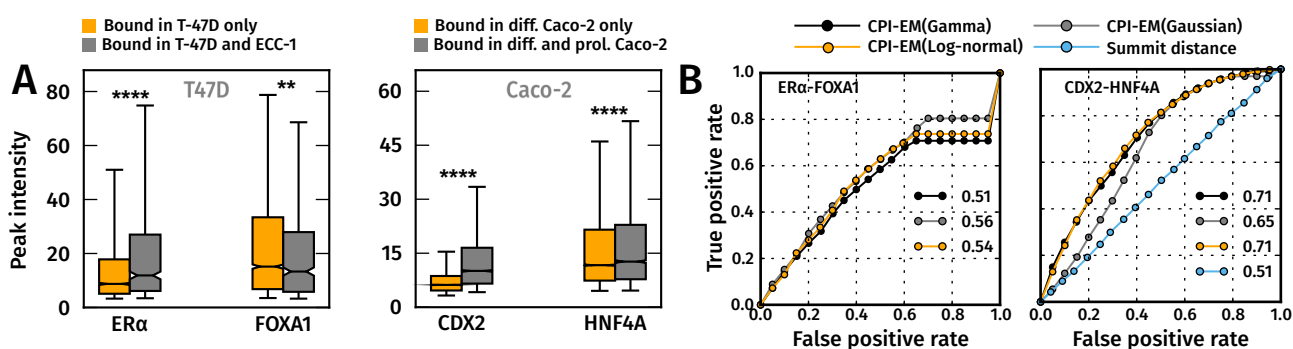


Figure 5: **(A) Regions bound by $ER\alpha$ only in T-47D cells are more weakly bound than regions bound by $ER\alpha$ in both T-47D and ECC-1 cells.** The same trend in peak intensities holds between regions bound by CDX2 only in differentiated Caco-2 cells and those bound by CDX2 in both differentiated and proliferating Caco-2 cells. However, cell-type specific binding in these cell types is also determined by factors other than cooperativity. Distributions of cooperatively and non-cooperatively bound regions are shown in orange and gray, respectively. The whiskers of the box plot are the 5 – *th* and 95 – *th* percentiles of the distributions shown. **(B) The Lognormal CPI-EM variant consistently detects cell-type specific binding events of $ER\alpha$ and CDX2.** ROC curves of Lognormal (orange), Gamma (black) and Gaussian (gray) variants of CPI-EM, and the peak distance detector (blue), on $ER\alpha$ -FOXA1 and CDX2-HNF4A datasets. The area under the ROC of each detector is indicated in the legend. The peak distance detector was not run on $ER\alpha$ -FOXA1 data since peak locations were not available in the peak calls. The ROC was generated using the same procedure as in Figure 4.