GimmeMotifs: an analysis framework for transcription factor motif analysis

This manuscript ([permalink](#)) was automatically generated from [simonvh/gimmemotifs-manuscript@6298f92](#) on November 20, 2018.

Authors

- **Niklas Bruse**

Radboud University, Faculty of Science, Department of Molecular Developmental Biology, Radboud Institute for Molecular Life Sciences, 6500 HB Nijmegen, The Netherlands

- **Simon J. van Heeringen**

[0000-0002-0411-3219](#) · [simonvh](#) · [svheeringen](#)

Radboud University, Faculty of Science, Department of Molecular Developmental Biology, Radboud Institute for Molecular Life Sciences, 6500 HB Nijmegen, The Netherlands · Funded by NWO-ALW, grant 863.12.002

Abstract

Background: Transcription factors (TFs) bind to specific DNA sequences, TF motifs, in cis-regulatory sequences and control the expression of the diverse transcriptional programs encoded in the genome. The concerted action of TFs within the chromatin context enables precise temporal and spatial expression patterns. To understand how TFs control gene expression it is essential to model TF binding. TF motif information can help to interpret the exact role of individual regulatory elements, for instance to predict the functional impact of non-coding variants.

Findings: Here we present GimmeMotifs, a comprehensive computational framework for TF motif analysis. Compared to the previously published version, this release adds a whole range of new functionality and analysis methods. It now includes tools for *de novo* motif discovery, motif scanning and sequence analysis, motif clustering, calculation of performance metrics and visualization. Included with GimmeMotifs is a non-redundant database of clustered motifs. Compared to other motif databases, this collection of motifs shows competitive performance in discriminating bound from unbound sequences. Using our *de novo* motif discovery pipeline we find large differences in performance between *de novo* motif finders on ChIP-seq data. Using an ensemble method such as implemented in GimmeMotifs will generally result in improved motif identification compared to a single motif finder. Finally, we demonstrate *maelstrom*, a new ensemble method that enables comparative analysis of TF motifs between multiple high-throughput sequencing experiments, such as ChIP-seq or ATAC-seq. Using a collection of ~200 H3K27ac ChIP-seq data sets we identify TFs that play a role in hematopoietic differentiation and lineage commitment.

Conclusion: GimmeMotifs is a fully-featured and flexible framework for TF motif analysis. It contains both command-line tools as well as a Python API and is freely available at: <https://github.com/vanheeringen-lab/gimmemotifs>.

Introduction

The regulatory networks that determine cell and tissue identity are robust, yet remarkably flexible. Transcription factors (TFs) control the expression of genes by binding to their cognate DNA sequences, TF motifs, in cis-regulatory elements [1]. To understand how genetic variation affects binding and to elucidate the role of TFs in regulatory networks we need to be able to accurately model binding of TFs to the DNA sequence.

The specificity of DNA-binding proteins can be modeled using various representations [2]. One of the most widely adopted is the position frequency matrix (PFM). This matrix, a TF motif, contains (normalized) frequencies of each nucleotide at each position in a collection of aligned binding sites. These PFMs can be derived from high-throughput experiments such as Chromatin Immunoprecipitation followed by sequencing (ChIP-seq) [3,4,5,6], HT-SELEX [7] or Protein Binding Microarrays (PBMs) [8]. Through straightforward transformations a PFM can be expressed as a weight matrix, using log likelihoods, or information content, using the Kullback-Leibler divergence.

We previously published GimmeMotifs, a *de novo* ChIP-seq motif discovery pipeline [9]. Here, we present a new and updated version of the GimmeMotifs software. Compared to the previous version, it now contains a whole range of Python modules and command-line tools to provide an comprehensive framework for transcription factor motif analysis. Amongst other possibilities it can be used to perform *de novo* motif analysis, cluster and visualize motifs and to calculate enrichment statistics. A new ensemble method, *maelstrom*, can be used to determine differential motif activity between multiple different conditions, such as cell types or treatments. We illustrate the functionality of GimmeMotifs using three different examples.

Findings

GimmeMotifs includes several different modules (Table 1), that can all be used either via a command-line tool or using the Python API. In the next sections we illustrate the functionality using three different examples.

Table 1: Command line tools included with GimmeMotifs. All these tools can also be used via the GimmeMotifs API.

Command	Purpose
<code>gimme motifs</code>	Find <i>de novo</i> motifs
<code>gimme maelstrom</code>	Identify differential motif activity
<code>gimme scan</code>	Scan for known motifs
<code>gimme roc</code>	Calculate performance metrics
<code>gimme match</code>	Find closest match in database
<code>gimme cluster</code>	Cluster motifs
<code>gimme background</code>	Generate background sequences
<code>gimme logo</code>	Create motif logo image

Benchmark of transcription factor motif databases

A variety of transcription factor (TF) motif databases have been published based on different data sources. One of the most established is JASPAR, which consists of a collection of non-redundant, curated binding profiles [10]. The JASPAR website contains many other tools and the underlying databases are also accessible via an API [11]. Other databases are based on protein binding micorarrays [8], HT-SELEX [7] or ChIP-seq profiles [3,4,5,6]. CIS-BP integrates many individual motif databases, and includes assignments of TFs to motifs bases based on DNA binding domain homology [12].

For the purpose of motif analysis, it is beneficial to have a database that is non-redundant (i.e., similar motifs are grouped together), yet as complete as possible (i.e., covers a wide variety of TFs). To establish a quantitative measure of database quality, we evaluated how well motifs from different databases can classify ChIP-seq peaks from background sequences. In this benchmark, we used randomly selected genomic sequence as background. The following databases were included in our comparison: JASPAR 2018 vertebrate [10], SwissRegulon [13], Homer [6], Factorbook [3], the ENCODE motifs from Kheradpour et al. [4], HOCOMOCO [5], the RSAT clustered motifs [14] and the IMAGE motif database created by Madsen et al. [15]. We compared these databases to the non-redundant vertebrate motif database included with GimmeMotifs (v5.0, see Methods).

Figure 1: Benchmark of transcription factor motif databases. A) Motif-based classification of binding sites for 294 TFs from the ReMap ChIP-seq database. For all TFs 5,000 peaks were compared to background regions using each motif database. The boxplot shows the the ROC AUC of the best motif per database for all TFs. Every point in this plot is based on one TF ChIP-seq peak set. **B)** Recall at 10% FDR of motif databases compared to the GimmeMotifs vertebrate motif database (v5.0). The same data is used as in **A)**. The X-axis represents the recall for the different databases, the Y-axis represents the recall for the GimmeMotifs vertebrate database. Differences of more than 0.025 are marked blue, and less then -0.025 red.

As a reference data set we downloaded all ChIP-seq peaks from ReMap 2018 [16], and selected the TFs with at least 1,000 peaks. When there were more than 5,000 peaks for a TF we randomly selected 5,000 peaks for the analysis.

We then evaluated the eight motif databases to test how well they could discriminate TF peaks from random genomic sequences using the GimmeMotifs tool `gimme roc`. This tool calculates a range of performance metrics to compare motif quality. For each TF and motif database combination we selected the single best performing motif, depending on the metric under consideration.

Figure 1A shows distribution of the ROC AUC (area under the curve for the Receiver Operator Curve) of the best motif per database for all 294 transcription factors in a box plot. The ROC curve plots the fraction of true positives (TPR, or sensitivity) against the fraction of false positives (FPR, or 1 - specificity). The ROC AUC value will generally range from 0.5 (no improvement on random guessing) to 1.0 (perfect classifier). As Figure 1A illustrates, the collection of TFs generally shows a wide distribution of ROC AUCs. For some factors, such as ELK1, CTCF, CBFB and MYOD1, peaks are relatively easy to classify using a single PFM motif. Other factors do not have peaks with a consistently enriched motif, or do not contain a sequence-specific DNA-binding domain, such as EP300 or CDK2 for example.

The difference in maximum ROC AUC between databases is on average not very large, with a mean maximum difference of 0.05. The largest difference (~0.24) is found for factors that were not assayed by ENCODE, such as ONECUT1, SIX2 and TP73, and are therefore not present in the Factorbook motif database. Unsurprisingly, the databases that were based on motif collections of different sources (ENCODE, IMAGE, RSAT and GimmeMotifs) generally perform best. It should be noted that, for this task, using motif databases based on motif identification from ChIP-seq peaks is in some sense “overfitting”, as the motifs in these databases were inferred from highly similar data.

While the ROC AUC is often used to compare the trade-off between sensitivity versus specificity, in this context it is not the best metric from a biological point of view. An alternative way of measuring performance is evaluating the recall (i.e. how many true peaks do we recover) at a specific false discovery rate. This is one of the criteria that has been used by the ENCODE DREAM challenge for evaluation [17]. Figure 1B shows scatterplots for the recall at 10% FDR for all motif databases compared to the clustered, non-redundant databases that is included with GimmeMotifs. This database shows better performance than most other databases using this benchmark. The non-redundant RSAT database, which was created in a very similar manner [14], scores comparably.

These results illustrate how `gimme roc` can be used for evaluation of motifs. The choice of a motif database can greatly influence the results of an analysis. The default database included with GimmeMotifs shows good performance on the metric evaluated here. However, this analysis illustrates only one specific use case of application of a motif database. In other cases well-curated databases such as JASPAR can be beneficial, for instance when linking motifs to binding proteins. Any of these databases can be easily used in all GimmeMotifs tools.

Large-scale benchmark of *de novo* motif finder performance on ChIP-seq peaks

It has been noted that there is no *de novo* motif prediction algorithm that consistently performs well across different data sets [18]. New approaches and algorithms for *de novo* motif discovery continue to be published, however, many of them are not tested on more than a few datasets. Benchmarks that have been published since Tompa et al. [19,20] typically have tested only a few motif finders or used only a few datasets.

Here, we used the GimmeMotifs framework as implemented in `gimme motifs` to benchmark 14 different *de novo* motif finders. To evaluate the different approaches, we downloaded 495 peak files for 270 proteins from ENCODE [21] and selected the 100bp sequence centered on the summit of top 5,000 peaks. These will be the peaks most likely to contain the primary TF motif, and should provide a straightforward test-case for the *de novo* motif finders. Ranking and selecting peaks in this manner is a widely adopted practice and we use this procedure also for our benchmark.

However, when analyzing ChIP-seq data in detail, it might be preferable to analyze to full complement of peaks.

Of the top peaks, half were randomly selected as a prediction set and the other half was used for evaluation. As a background set we selected regions of the same length flanking the original peaks. This will account for sequence bias according to genomic distribution. To assess the performance, we calculated two metrics, the ROC AUC and the recall at 10% FDR. Figure 2A shows the distribution of the ROC AUC scores over all ENCODE peaks in a boxplot, ordered by the mean ROC AUC. The ROC AUC is distributed between 0.58 and 0.98, with a mean of 0.75. All proteins that have low ROC AUC are not sequence-specific transcription factors such as POL2, TAF7 and GTF2B, the PRC2-subunit SUZ12 and the H3K9 methyltransferase SETDB1. The factors with the highest ROC AUC are CTCF and members of the cohesin complex, SMC3 and RAD21, that bind at CTCF sites.

Figure 2: Benchmark of de novo motif finders. A) Comparison of the ROC AUC of the best motif of each motif finder. The boxplot shows the best motif per peak set of 495 peaks for 270 proteins from ENCODE. The best motif from all motif finders is indicated as 'Best'. **B)** Comparison of the best motif per motif finder compared to the best overall motif for each data set. Plotted is the difference in recall compared to the best motif. Recall is calculated at 10% FDR. **C)** The relative motif rank as a function of the motif quality. Rank is the mean overall rank of three metrics (ROC AUC, recall at 10% FDR and MNCP).

Generally, the ROC AUC distribution of all evaluated motif finders is very similar. However, a few outliers can be observed. Trawler and Posmo show an overall lower distribution of ROC AUC scores. Compared to the ROC AUC scores of the next best program, GADeM, this is significant ($p < 0.01$, Wilcoxon signed-rank). Selecting the best motif for each experiment, as the ensemble method implemented in GimmeMotifs would do, results in a ROC AUC distribution that is significantly higher than the best single method, BioProspector ($p < 1e-21$, Wilcoxon signed-rank).

As stated in the previous section, the ROC AUC is not the best measure to evaluate motif quality. Therefore, we selected for every TF peak set the best overall motif on the basis of the recall at 10%

FDR. We then plotted the difference between this best overall motif and the best motif from each individual *de novo* approach (Fig. 2B). For this figure, we used only the data sets where at least one motif had a recall higher than 0 at 10% FDR.

In line with previous results [18], there is no single tool that consistently predicts the best motif for each transcription factor. However, the motifs predicted by BioProspector, MEME and Homer are, on basis of this metric, consistently better than motifs predicted by other methods. In 75% of the cases, the motif predicted by BioProspector has a difference in recall smaller than 0.026 compared to the best overall motif. In this benchmark, four programs (Trawler, Improbizer, Posmo and Weeder) generally perform worse than average, with a mean decrease in recall of 0.11 to 0.17, as compared to the best motif. In addition, these programs tend to have a much more variable performance overall.

Predicted motifs identified using MEME with different motif widths show better performance than running MEME with the `minw` and `maxw` options (MEME vs. MEMEW in Fig. 2B). Of the best performing algorithms, both MEME and BioProspector were not specifically developed for ChIP-seq data, however, they consistently outperform most methods created for ChIP-seq data. Of the ChIP-seq motif finders Homer consistently shows good performance.

Finally, to gain further insight into *de novo* motif finder performance, we stratified the ChIP-seq datasets by motif “quality”. We divided the transcription factors into five bins on basis of the ROC AUC score of the best motif. For each bin we ranked the tools on basis of the average of three metrics (ROC AUC, recall at 10% FDR and MNCP [22]). The results are visualized as a heatmap in Figure 2C. From this visualization, it is again clear that BioProspector, MEME and Homer produce consistently high-ranking motifs, while the motifs identified by Trawler, Posmo, GADeM and Improbizer generally have the lowest rank. Interestingly, for some motif finders, there is a relation between motif presence and the relative rank. Weeder, XXMotif and MDmodule yield relatively high-ranking motifs when the ROC AUC of the best motif for the data set is low. On the other hand, ChIPMunk shows the opposite pattern. Apparently this algorithm works well when a motif is present in a significant fraction of the data set.

These results illustrate that motif finders need to be evaluated along a broad range of data sets with different motif presence and quality. Another interesting observation is that this ChIP-seq benchmark shows a lower-than-average performance for Weeder, which actually was one of the highest scoring in the Tompa et al. benchmark. It should be noted that our metric specifically evaluates how well *de novo* motif finders identify the primary motif in the context of high-scoring ChIP-seq peaks. It does not evaluate other aspects that might be important, such as the ability to identify low-abundant or co-factor motifs. Furthermore, with ChIP-seq data there are usually thousands of peaks available. This allows for other algorithms than those that work well on a few sequences. Interestingly, the original MEME shows consistently good performance, although the running time is longer than most other tools. On the basis of this analysis, BioProspector should be

the top pick for a program to identify primary motifs in ChIP-seq data. However, an ensemble program such as GimmeMotifs will report high-quality motifs more consistently than any single tool.

Differential motif analysis of hematopoietic enhancers identifies cell type-specific regulators

While many motif scanners and methods to calculate enrichment exist, there are few methods to compare motif enrichment or activity between two or more data sets. The CentriMo algorithm from the MEME suite implements a differential enrichment method to compare two samples [23]. Other approaches, such as MARA [24,25] and IMAGE [15], are based on linear regression. Here we present the *maelstrom* algorithm that integrates different methods to determine motif relevance or activity in an ensemble approach (Fig. 3A).

To demonstrate the utility of *maelstrom* we identified motif activity based on enhancers in hematopoietic cells. We downloaded 69 human hematopoietic DNaseI experiments (Supplementary Table S1), called peaks, and created a combined peak set as a collection of putative enhancers. In addition we downloaded 193 hematopoietic H3K27ac ChIP-seq experiments, mainly from BLUEPRINT [26] (Supplementary Table S1). We determined the number of H3K27ac reads per enhancer (Supplementary Table S2). After log2 transformation and scaling, we selected the 50,000 most dynamic peaks. Figure 3B shows the correlation of the H3K27ac enrichment in these 50,000 enhancers between cell types. For this plot, replicates were combined by taking the mean value and all experiments corresponding to treated cells were removed. We can observe six main clusters 1) non-hematopoietic cells 2) neutrophilic cells, 3) monocytes, macrophages and dendritic cells, 4) megakaryocytes and erythroblasts, 5) B cells 6) T cells and natural killer (NK) cells. We can conclude that the H3K27ac profile within this enhancer set recapitulates a cell type-specific regulatory signal.

Figure 3: Predicting TF motif activity using maelstrom. **A)** A schematic overview of the maelstrom ensemble method. **B)** Heatmap of the correlation of H3K27ac signal in hematopoietic enhancers. We counted H3K27ac ChIP-seq reads in 2kb sequences centered at DNase I peaks. Counts were log2-transformed and scaled and replicates were combined by taking the mean value. This heatmap shows the Pearson r , calculated using the 50,000 most dynamic peaks. **C)** Selection of the results of running `gimme maelstrom` on the 50,000 most dynamic hematopoietic enhancers. Shown is the motif activity of four motifs associated with factors that are known to play a role in

*hematopoietic cells: SPI1 (PU.1), CEBP, RUNX and GATA1. The visualization shows a schematic, simplified cell lineage tree. The color and the line thickness represent the motif activity, where the value corresponds to the $\log_{10}$ of the p-value of the rank aggregation. For high-ranking motifs (red) $-\log_{10}(p\text{-value})$ is shown, while for low ranking motifs (blue) $\log_{10}(p\text{-value})$ of the reversed ranking is shown. **D)** Motif activity, as in **C**, of two motifs of factors for which the exact function in these cell types is currently unknown.*

To determine differential motif activity from these dynamic enhancers we used maelstrom. We combined Lasso, Bayesian ridge regression, multi-class regression using coordinate descent [27] and regression with boosted trees [28]. The coefficients or feature importances were ranked and combined using rank aggregation [29]. A p-value was calculated for consistently high ranking and consistently low ranking motifs. A selection of the results is visualized in Figure 3C. The full results are available as Supplementary File S1 and on Zenodo ([10.5281/zenodo.1491482](https://doi.org/10.5281/zenodo.1491482)).

Two of the most significant motifs are SPI1 (PU.1) and CEBP (Fig. 3C). The motif activity for SPI1 is high in monocytes and macrophages, consistent with its role in myeloid lineage commitment [30]. The CEBP family members are important for monocytes and granulocytic cells [31], and show a high motif activity in neutrophils and monocytes. Other strong motifs include RUNX for T cells and NK cells and GATA1 for erythroid cells (Fig. 3C).

We identified a high activity for motifs representing the ZEB1 and Snail transcription factors (Fig. 3D). The Snail transcription factors play an important role in the epithelial-to-mesenchymal transition (EMT), and their role in hematopoietic cells is less well-described. However, recently Snai2 and Snai3 were found to be required to generate mature T and B cells [32,33] in mice. ZEB1 is expressed in T cells and represses expression of IL-2 [34], as well as other immune genes such as CD4 [35] and GATA3 [36]. ZEB1 knockout mice exhibit a defect in thymocyte development [37]. Together, this suggests that these TFs could play an important role in the hematopoietic lineage.

Finally, an interesting observation is the predicted motif activity of NANOG in endothelial cells (Fig. 3D). NANOG is expressed in embryonic stem cells and is essential for maintenance of pluripotency [38]. However, NANOG is indeed also expressed in endothelial cells and has been shown to play a role in endothelial proliferation and angiogenesis [39].

Figure 4: Predicting TF motif activity using *maelstrom*. Results of running `gimme maelstrom` on the 50,000 most dynamic hematopoietic enhancers. The motif activity of the top 78 motifs (absolute motif activity ≥ 5) is visualized in a heatmap. The color represents the reported motif activity, where the value corresponds to the $\log_{10}$ of the p -value of the rank aggregation. For high-ranking motifs (red) $-\log_{10}(p\text{-value})$ is shown, while for low ranking motifs (blue) $\log_{10}(p\text{-value})$ of the reversed ranking is shown.

In addition to the described examples, we identified a large compendium of TF motifs that display differential activity in the hematopoietic lineage (Fig. 4). This demonstrates that `gimme maelstrom` can be used to analyze complex, multi-dimensional data sets such as this large-scale collection of hematopoietic enhancers. Especially in experiments where there are multiple conditions or time points that need to be compared, `maelstrom` provides a powerful method to determine differential transcription factor motif activity.

Methods

GimmeMotifs

Implementation

GimmeMotifs is implemented in Python, with the motif scanning incorporated as a C module. The software is developed on GitHub (<https://github.com/vanheeringen-lab/gimmemotifs/>) and documentation is available at <https://gimmemotifs.readthedocs.io>. Functionality is covered by unit tests, which are run through continuous integration. GimmeMotifs can be installed via bioconda [40], see <https://bioconda.github.io/> for details. All releases are also distributed through PyPi [41] and stably archived using Zenodo [42]. For *de novo* motif search, 14 different external tools are supported (Table 2). All of these are installed when conda is used for installation. By default, `genomepy` is used for genome management [43]. In addition, GimmeMotifs uses the following Python modules: numpy [44], scipy [45], scikit-learn, scikit-contrib-lightning [27], seaborn [46], pysam [47,48], xgboost [28] and pandas. In addition to the command line tools, all GimmeMotifs functionality is available through a Python API.

De novo motif prediction pipeline

Originally, GimmeMotifs was developed to predict *de novo* motifs from ChIP-seq data using an ensemble of motif predictors [9]. The tools currently supported are listed in Table 2. An input file (BED, FASTA or narrowPeak format) is split into a prediction and validation set. The prediction set is used to predict motifs, and the validation set is used to filter for significant motifs. All significant motifs are clustered to provide a collection of non-redundant *de novo* motifs. Finally, significant clustered motifs are reported, along with several statistics to evaluate motif quality, calculated using

the validation set. These evaluation metrics include ROC AUC, distribution of the motif location relative to the center of the input (i.e., the ChIP-seq peak summit) and the best match in a database of known motifs.

Table 2: External *de novo* motif prediction tools supported by GimmeMotifs. The version shown is the version string reported by the software. If the version is unknown, the date when the tool was added to the GimmeMotifs repository is shown.

Name	Version	Citation
AMD	Unknown (2012-11-21)	[49]
BioProspector	4/15/04	[50]
ChIPMunk	V7 10012017	[51]
GADEM	v1.3.1	[52]
HMS	Unknown (2012-11-23)	[53]
Homer	2	[6]
Improbizer	Unknown (2010-10-01)	[54]
MDmodule	Unknown (2010-07-29)	[55]
MEME	4.6.0	[56]
MotifSampler	3.2	[57]
Posmo	May-19-2011	[58]
Trawler	2.0	[59]
Weeder	2.0	[60]
XXmotif	1.6	[61]

Motif activity by ensemble learning: maelstrom

GimmeMotifs implements eight different methods to determine differential enrichment of known motifs between two or more conditions. In addition, these methods can be combined in a single measure of *motif activity* using rank aggregation. Four methods work with discrete sets, such as different peak sets or clusters from a K-means clustering. The hypergeometric test uses motif counts with an empirical motif-specific FPR of 5%. All other implemented methods use the z-score normalized PFM log-odds score of the best match.

To combine different measures of motif significance or activity into a single score, ranks are assigned for each individual method and combined using rank aggregation based on order statistics [29]. This results in a probability of finding a motif at all observed positions. We use a Python implementation based on the method used in the R package RobustRankAggreg [62]. The rank aggregation is performed twice, once with the ranks reversed to generate both positively and negatively associated motifs.

The hypergeometric test is commonly used to calculate motif enrichment, for instance by Homer [6]. In GimmeMotifs, motifs in each cluster are tested against the union of all other clusters. The reported value is $-\log_{10}(\text{p-value})$ where the p-value is adjusted by the Benjamini-Hochberg procedure [63].

Using the non-parametric Mann-Whitney U test, GimmeMotifs tests the null hypothesis that the motif log-odds score distributions of two classes are equal. For each discrete class in the data, such as a cluster, it compares the score distributions of the class to the score distribution of all other classes. The value used as activity is the $-\log_{10}$ of the Benjamini-Hochberg adjusted p-value.

The two other methods are classification algorithms: random forest using scikit-learn and a large-scale multiclass classifier using block coordinate descent [27] as implemented in the scikit-contrib-lightning module. The classifier in GimmeMotifs uses a l1/l2 penalty with squared hinge loss where the alpha and C parameters are set using grid search in 10 fold cross-validation procedure.

The other four methods that are implemented relate motif score to an experimental measure such as ChIP-seq or ATAC-seq signal or expression level. These are all different forms of regression. In addition to ridge regression, which is similar to Motif Activity Response Analysis (MARA) [24,25], these methods include regression using boosted trees (XGBoost [28]), multiclass regression [27] and L1 regularized regression (LASSO).

Clustering to create the gimme.vertebrate.v5.0 motif database

We collected all motifs from the following databases: CIS-BP v1.02 [12], ENCODE [4], Factorbook [3], HOCOMOCO v11 [5], HOMER v4.10 [6], JASPAR 2018 vertebrates [10] and SwissRegulon [13]. Motif similarity was calculated using Pearson correlation of motif scores profiles [64,65] using a sequence that contains each 7-mer or its reverse complement [66]. We then clustered the motifs using agglomerative clustering with complete linkage and connectivity constraints where only motifs with a Pearson $r \geq 0.5$ were considered as neighbors. The number of clusters was set to 1900. After clustering, we discarded all motifs where the sum of the information content of all positions was less than 5. The clustered database, `gimme.vertebrate.v5.0`, is distributed with GimmeMotifs.

Transcription factor motif database benchmark

We downloaded all hg38 non-ENCODE ChIP-seq peaks from Remap 2018 v1.2 [16] (<http://tagc.univ-mrs.fr/remap/index.php?page=download>). We removed all factors with fewer than 1000 peaks and created regions of 100 bp centered at the peak summit. Background files were created for each peak set using bedtools shuffle [67], excluding the hg38 gaps and the peak regions. The ROC AUC and Recall at 10% FDR statistics were calculated using `gimme roc`. The motif databases included in the comparison are listed in Table 3. We only included public databases that can be freely accessed and downloaded.

Table 3: Motif databases.

Name	Version	Citation
Factorbook	Sep. 2012	[3]
SwissRegulon	Nov. 2006	[13]
GimmeMotifs vertebrate	v5.0	
HOCOMOCO	v11	[5]
Homer	v4.10	[6]
JASPAR	2018	[10]
ENCODE	Dec. 2013	[4]
Madsen	1.1	[15]
RSAT vertebrate clusters	Sep. 2017	[14]

The workflow is implemented in snakemake [68] and is available at <https://github.com/vanheeringen-lab/gimme-analysis>.

`De novo` motif prediction benchmark

We downloaded all spp ENCODE peaks (January 2011 data freeze) from the EBI FTP (http://ftp.ebi.ac.uk/pub/databases/ensembl/encode/integration_data_jan2011/byDataType/peaks/jan2011/spp/optimal/). We selected the top 5000 peaks and created 100bp regions centered on the peak summit. As background we selected 100 bp regions flanking the original peaks. For the `de novo` motif search default settings for `gimme motifs` were used. The workflow is implemented in snakemake [68] and is available at <https://github.com/vanheeringen-lab/gimme-analysis>.

Motif analysis of hematopoietic enhancers

To illustrate the functionality of `gimme maelstrom` we analyzed an integrated collection of hematopoietic enhancers. We downloaded all H3K27ac ChIP-seq and DNase I data from BLUEPRINT and hematopoietic DNase I data from ROADMAP (Supplementary Table S1). All DNase I data were processed using the Kundaje lab DNase pipeline version 0.3.0 https://github.com/kundajelab/atac_dnase_pipelines [69]. The ChIP-seq samples were processed using the Kundaje lab AQUAS TF and histone ChIP-seq pipeline https://github.com/kundajelab/chipseq_pipeline. For all experiments from BLUEPRINT we used the aligned reads provided by EBI. All ROADMAP samples were aligned using bowtie2 [70] to the hg38 genome. DNase I peaks were called using MACS2 [71] according to the default settings of the DNase pipeline. We merged all DNase I peak files and centered each merged peak on the summit of the strongest individual peak. H3K27ac reads were counted in a region of 2kb centered at the summit (Supplementary Table S2) and read counts were log2-transformed and scaled. We removed all samples that were treated and averaged all samples from the same cell type. We then selected all enhancers with at least one sample with a scaled log2 read count of 2, sorted by the maximum difference in normalized signal between samples and selected the 50,000 enhancers with the largest difference. Using this enhancer collection as input, we ran `gimme maelstrom` using default settings. The motif analysis workflow is implemented in a Jupyter notebook and is available at <https://github.com/vanheeringen-lab/gimme-analysis>.

Conclusions

We demonstrated the functionality of GimmeMotifs with three examples. First, to evaluate different public motif databases, we quantified their performance on distinguishing ChIP-seq peaks from background sequences. The databases that perform best on this benchmark are collections of motifs from different sources. Of the individual databases HOCOMOCO and Factorbook rank highest using this collection of human ChIP-seq peaks. Based on our results it is recommended to use a composite database, such as the RSAT clustered motifs or the GimmeMotifs database (v5.0), for the best vertebrate motif coverage. However, these motifs are less well annotated. For instance, motifs based on ChIP-seq peaks from some sources might be from co-factors or cell type-specific regulators instead of the factor that was assayed. An example are motifs that are associated with the histone acetyl transferase EP300. This transcriptional co-activator lacks a DNA binding domain, and associated motifs depend on the cell type. For instance, in a lymphoblastoid cell line such as GM12878 these include PU.1 and AP1. The lack of high-quality annotation makes it more difficult to reliably link motifs to transcription factors. This can be a distinct advantage of a database such as JASPAR. Although the motifs might not be optimal for every TF, JASPAR contains high-quality metadata that is manually curated.

In the second example, we benchmarked 14 different *de novo* motif finders using a large compendium of ChIP-seq data. While performance can vary between different data sets, there are several *de novo* motif finders that consistently perform well, with BioProspector, MEME and Homer as top performers. Interestingly, only Homer was specifically developed for ChIP-seq data. An ensemble approach such as GimmeMotifs still improves on the use of individual tools. This example also illustrates that newly developed *de novo* motif finders should be evaluated on many different data sets, as this is necessary to accurately judge the performance.

Finally, we presented a new ensemble approach, *maelstrom*, to determine motif activity in two or more epigenomic or transcriptomic data sets. Using H3K27ac ChIP-seq signal as a measure for enhancer activity, we analyzed cell-type specific motif activity in a large collection of hematopoietic cell types. We identified known lineage regulators, as well as motifs for factors that are less well studied in a hematopoietic context. This illustrates how `gimme maelstrom` can serve to identify cell type-specific transcription factors and has the potential to discover

In conclusion, GimmeMotifs is a flexible and highly versatile framework for transcription factor motif analysis. Both command line and programmatic use in Python are supported. One planned future improvement to GimmeMotifs is the support of more sophisticated motif models. Even though the PFM is a convenient representation, it has certain limitations. A PFM cannot model inter-nucleotide dependencies, which are known to affect binding of certain TFs. Multiple different representations have been proposed [72,73,74,75,76], but no single one of these has gained much traction. It is still unclear how well these models perform and their use depends on specific tools. Supporting these different models and benchmarking their performance relative to high-quality PFMs will simplify their use and give insight into their benefits and disadvantages. Second, there is significant progress recently in modeling TF binding using deep neural networks (DNNs) [77,78]. These DNNs can learn sequence motifs, as well as complex inter-dependencies, directly from the data. However, while biological interpretation is possible [79], it becomes less straightforward. We expect that analyzing and understanding a trained DNN can benefit from high-quality motif databases and comparative tools such as GimmeMotifs.

Availability and requirements

- Project name: GimmeMotifs
- Project home page: <https://github.com/vanheeringen-lab/gimmemotifs>
- Operating system(s): Linux, Mac OSX
- Programming language: Python 3
- Other requirements: *de novo* motif finders
- License: MIT

Availability of supporting data

- Scripts and notebooks to reproduce the analysis are available at:
 - <https://github.com/vanheeringen-lab/gimme-analysis>.
- Additional data files are available at Zenodo:
 - Table of H3K27ac read counts at all DNase I accessible sites: [10.5281/zenodo.1488669](https://zenodo.org/record/1488669).
 - Results of gimme maelstrom (Figure 3 and 4): [10.5281/zenodo.1491482](https://zenodo.org/record/1491482).

Additional files

- Supplementary Table S1: List of accessions used in the analysis of hematopoietic enhancers.
- Supplementary Table S2: Table of H3K27ac read counts at all DNase I accessible sites.
- Supplementary File S1: Results of gimme maelstrom (Figure 3 and 4).

Competing interests

The authors declare that they have no competing interests.

Funding

SJvH was supported by the Netherlands Organization for Scientific research (NWO-ALW, grants 863.12.002 and 016.Vidi.189.081). Part of this work was carried out on the Dutch national e-infrastructure with the support of SURF Foundation. This work was sponsored by NWO Exact and Natural Sciences for the use of supercomputer facilities.

Author contributions

SJvH designed and performed experiments and analysis, wrote the code and wrote the manuscript. NB processed and analyzed the hematopoietic DNase I and H3K27ac ChIP-seq data. All authors read and approved the final manuscript.

Acknowledgements

We would like to thank all research symbionts, who make their tools and data publicly available. This study makes use of data generated by the Blueprint Consortium. A full list of the investigators who contributed to the generation of the data is available from www.blueprint-epigenome.eu.

Additionally, this study used data provided by the NIH Roadmap Epigenomics Consortium (<http://nihroadmap.nih.gov/epigenomics/>).

References

1. The Human Transcription Factors

Samuel A. Lambert, Arttu Jolma, Laura F. Campitelli, Pratyush K. Das, Yimeng Yin, Mihai Albu, Xiaoting Chen, Jussi Taipale, Timothy R. Hughes, Matthew T. Weirauch
Cell (2018-02) <https://doi.org/10.1016/j.cell.2018.01.029>

2. Modeling the specificity of protein-DNA interactions

Gary D. Stormo
Quantitative Biology (2013-04-02) <https://doi.org/10.1007/s40484-013-0012-4>

3. Factorbook.org: a Wiki-based database for transcription factor-binding data generated by the ENCODE consortium

J. Wang, J. Zhuang, S. Iyer, X.-Y. Lin, M. C. Greven, B.-H. Kim, J. Moore, B. G. Pierce, X. Dong, D. Virgil, ... Z. Weng
Nucleic Acids Research (2012-11-29) <https://doi.org/10.1093/nar/gks1221>

4. Systematic discovery and characterization of regulatory motifs in ENCODE TF binding experiments

P. Kheradpour, M. Kellis
Nucleic Acids Research (2013-12-13) <https://doi.org/10.1093/nar/gkt1249>

5. HOCOMOCO: towards a complete collection of transcription factor binding models for human and mouse via large-scale ChIP-Seq analysis

Ivan V Kulakovskiy, Ilya E Vorontsov, Ivan S Yevshin, Ruslan N Sharipov, Alla D Fedorova, Eugene I Rumynskiy, Yulia A Medvedeva, Arturo Magana-Mora, Vladimir B Bajic, Dmitry A Papatsenko, ... Vsevolod J Makeev
Nucleic Acids Research (2017-11-11) <https://doi.org/10.1093/nar/gkx1106>

6. Simple Combinations of Lineage-Determining Transcription Factors Prime cis-Regulatory Elements Required for Macrophage and B Cell Identities

Sven Heinz, Christopher Benner, Nathanael Spann, Eric Bertolino, Yin C. Lin, Peter Laslo, Jason X. Cheng, Cornelis Murre, Harinder Singh, Christopher K. Glass
Molecular Cell (2010-05) <https://doi.org/10.1016/j.molcel.2010.05.004>

7. Multiplexed massively parallel SELEX for characterization of human transcription factor binding specificities

A. Jolma, T. Kivioja, J. Taipale, L. Cheng, G. Wei, M. Enge, M. Taipale, J. M. Vaquerizas, J. Yan, M. J. Sillanpaa, ... J. Taipale
Genome Research (2010-04-08) <https://doi.org/10.1101/gr.100552.109>

8. UniPROBE, update 2015: new tools and content for the online database of protein-binding microarray data on protein–DNA interactions

Maxwell A. Hume, Luis A. Barrera, Stephen S. Gisselbrecht, Martha L. Bulyk

Nucleic Acids Research (2014-11-05) <https://doi.org/10.1093/nar/gku1045>

9. GimmeMotifs: a de novo motif prediction pipeline for ChIP-sequencing experiments

Simon J. van Heeringen, Gert Jan C. Veenstra

Bioinformatics (2010-11-15) <https://doi.org/10.1093/bioinformatics/btq636>

10. JASPAR 2018: update of the open-access database of transcription factor binding profiles and its web framework

Aziz Khan, Oriol Fornes, Arnaud Stigliani, Marius Gheorghe, Jaime A Castro-Mondragon, Robin van der Lee, Adrien Bessy, Jeanne Chèneby, Shubhada R Kulkarni, Ge Tan, ... Anthony Mathelier
Nucleic Acids Research (2017-11-13) <https://doi.org/10.1093/nar/gkx1126>

11. JASPAR RESTful API: accessing JASPAR data from any programming language

Aziz Khan, Anthony Mathelier

Bioinformatics (2017-12-15) <https://doi.org/10.1093/bioinformatics/btx804>

12. Determination and Inference of Eukaryotic Transcription Factor Sequence Specificity

Matthew T. Weirauch, Ally Yang, Mihai Albu, Atina G. Cote, Alejandro Montenegro-Montero, Philipp Drewe, Hamed S. Najafabadi, Samuel A. Lambert, Ishminder Mann, Kate Cook, ... Timothy R. Hughes

Cell (2014-09) <https://doi.org/10.1016/j.cell.2014.08.009>

13. SwissRegulon: a database of genome-wide annotations of regulatory sites

M. Pachkov, I. Erb, N. Molina, E. van Nimwegen

Nucleic Acids Research (2007-01-03) <https://doi.org/10.1093/nar/gkl857>

14. RSAT matrix-clustering: dynamic exploration and redundancy reduction of transcription factor binding motif collections

Jaime Abraham Castro-Mondragon, Sébastien Jaeger, Denis Thieffry, Morgane Thomas-Chollier, Jacques van Helden

Nucleic Acids Research (2017-06-07) <https://doi.org/10.1093/nar/gkx314>

15. Integrated analysis of motif activity and gene expression changes of transcription factors

Jesper Grud Skat Madsen, Alexander Rauch, Elvira Laila Van Hauwaert, Søren Fisker Schmidt, Marc Winnefeld, Susanne Mandrup

Genome Research (2017-12-12) <https://doi.org/10.1101/gr.227231.117>

16. ReMap 2018: an updated atlas of regulatory regions from an integrative analysis of DNA-binding ChIP-seq experiments

Jeanne Chèneby, Marius Gheorghe, Marie Artufel, Anthony Mathelier, Benoit Ballester

Nucleic Acids Research (2017-11-08) <https://doi.org/10.1093/nar/gkx1092>

17. ENCODE-DREAM in vivo Transcription Factor Binding Site Prediction Challenge - Dream Challenges
Dream Challenges <http://dreamchallenges.org/project/encode-dream-in-vivo-transcription-factor-binding-site-prediction-challenge/>

18. Assessing computational tools for the discovery of transcription factor binding sites

Martin Tompa, Nan Li, Timothy L Bailey, George M Church, Bart De Moor, Eleazar Eskin, Alexander V Favorov, Martin C Frith, Yutao Fu, W James Kent, ... Zhou Zhu

Nature Biotechnology (2005-01) <https://doi.org/10.1038/nbt1053>

19. Evaluating tools for transcription factor binding site prediction

Narayan Jayaram, Daniel Usvyat, Andrew C. R. Martin

BMC Bioinformatics (2016-11-02) <https://doi.org/10.1186/s12859-016-1298-9>

20. Limitations and potentials of current motif discovery algorithms

J. Hu, B. Li, D. Kihara

Nucleic Acids Research (2005-09-02) <https://doi.org/10.1093/nar/gki791>

21. Index of /pub/databases/ensembl/encode/integration_data_jan2011/byDataType/peaks/jan2011/spp/optimal/hub
http://ftp.ebi.ac.uk/pub/databases/ensembl/encode/integration_data_jan2011/byDataType/peaks/jan2011/spp/optimal/hub/

22. Rank order metrics for quantifying the association of sequence features with gene regulation.

Neil D Clarke, Joshua A Granek

Bioinformatics (Oxford, England) (2003-01-22) <https://www.ncbi.nlm.nih.gov/pubmed/12538241>

23. Differential motif enrichment analysis of paired ChIP-seq experiments

Tom Lesluyes, James Johnson, Philip Machanick, Timothy L Bailey

BMC Genomics (2014) <https://doi.org/10.1186/1471-2164-15-752>

24. The transcriptional network that controls growth arrest and differentiation in a human myeloid leukemia cell line

Harukazu Suzuki, Alistair RR Forrest, Erik van Nimwegen, Carsten O Daub, Piotr J Balwierz, Katharine M Irvine, Timo Lassmann, Timothy Ravasi, Yuki Hasegawa, ...

Nature Genetics (2009-04-19) <https://doi.org/10.1038/ng.375>

25. ISMARA: automated modeling of genomic signals as a democracy of regulatory motifs

P. J. Balwierz, M. Pachkov, P. Arnold, A. J. Gruber, M. Zavolan, E. van Nimwegen

Genome Research (2014-02-10) <https://doi.org/10.1101/gr.169508.113>

26. BLUEPRINT: mapping human blood cell epigenomes

J. H. A. Martens, H. G. Stunnenberg

Haematologica (2013-10-01) <https://doi.org/10.3324/haematol.2013.094243>

27. Block coordinate descent algorithms for large-scale sparse multiclass classification

Mathieu Blondel, Kazuhiro Seki, Kuniaki Uehara

Machine Learning (2013-05-08) <https://doi.org/10.1007/s10994-013-5367-2>

28. XGBoost

Tianqi Chen, Carlos Guestrin

Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16 (2016) <https://doi.org/10.1145/2939672.2939785>

29. Gene prioritization through genomic data fusion

Stein Aerts, Diether Lambrechts, Sunit Maity, Peter Van Loo, Bert Coessens, Frederik De Smet, Leon-Charles Tranchevent, Bart De Moor, Peter Marynen, Bassem Hassan, ... Yves Moreau

Nature Biotechnology (2006-05) <https://doi.org/10.1038/nbt1203>

30. PU.1 induces myeloid lineage commitment in multipotent hematopoietic progenitors.

C Nerlov, T Graf

Genes & development (1998-08-01) <https://www.ncbi.nlm.nih.gov/pubmed/9694804>

31. Transcriptional control of granulocyte and monocyte development

AD Friedman

Oncogene (2007-10) <https://doi.org/10.1038/sj.onc.1210764>

32. Fatal autoimmunity results from the conditional deletion of Snai2 and Snai3

Peter D. Pioli, Xinjian Chen, Janis J. Weis, John H. Weis

Cellular Immunology (2015-05) <https://doi.org/10.1016/j.cellimm.2015.02.009>

33. Snai2 and Snai3 transcriptionally regulate cellular fitness and functionality of T cell lineages through distinct gene programs

Peter D. Pioli, Sarah K. Whiteside, Janis J. Weis, John H. Weis

Immunobiology (2016-05) <https://doi.org/10.1016/j.imbio.2016.01.007>

34. The transcription repressor, ZEB1, cooperates with CtBP2 and HDAC1 to suppress IL-2 gene activation in T cells

J. Wang, S. Lee, C. E.-Y. Teh, K. Bunting, L. Ma, M. F. Shannon

International Immunology (2009-01-15) <https://doi.org/10.1093/intimm/dxn143>

35. Negative regulation of CD4 expression in T cells by the transcriptional repressor ZEB.

T Brabletz, A Jung, F Hlubek, C Löhberg, J Meiler, U Suchy, T Kirchner

International immunology (1999-10) <https://www.ncbi.nlm.nih.gov/pubmed/10508188>

36. T-cell expression of the human GATA-3 gene is regulated by a non-lineage-specific silencer.

JM Grégoire, PH Roméo

The Journal of biological chemistry (1999-03-05) <https://www.ncbi.nlm.nih.gov/pubmed/10037751>

37. DeltaEF1, a zinc finger and homeodomain transcription factor, is required for skeleton patterning in multiple lineages.

T Takagi, H Moribe, H Kondoh, Y Higashi

Development (Cambridge, England) (1998-01) <https://www.ncbi.nlm.nih.gov/pubmed/9389660>

38. Nanog and transcriptional networks in embryonic stem cell pluripotency

Guangjin Pan, James A Thomson

Cell Research (2007-01) <https://doi.org/10.1038/sj.cr.7310125>

39. NANOG induction of fetal liver kinase-1 (FLK1) transcription regulates endothelial cell proliferation and angiogenesis

E. E. Kohler, C. E. Cowan, I. Chatterjee, A. B. Malik, K. K. Wary

Blood (2010-11-30) <https://doi.org/10.1182/blood-2010-07-295261>

40. Bioconda: sustainable and comprehensive software distribution for the life sciences

Björn Grüning, Ryan Dale, Andreas Sjödin, Brad A. Chapman, Jillian Rowe, Christopher H. Tomkins-Tinch, Renan Valieris, Johannes Köster

Nature Methods (2018-07) <https://doi.org/10.1038/s41592-018-0046-7>

41. PyPI – the Python Package Index [PyPI https://pypi.org/](https://pypi.org/)

42. Zenodo - Research. Shared. <https://zenodo.org/>

43. genomepy: download genomes the easy way

Simon J. van Heeringen

The Journal of Open Source Software (2017-08-15) <https://doi.org/10.21105/joss.00320>

44. NumPy — NumPy(2018-04-26) <http://www.numpy.org/>

45. SciPy.org — SciPy.org(2018-08-15) <http://www.scipy.org/>

46. Mwaskom/Seaborn: V0.9.0 (July 2018)

Michael Waskom, Olga Botvinnik, Drew O'Kane, Paul Hobson, Joel Ostblom, Saulius Lukauskas, David C Gemperline, Tom Augspurger, Yaroslav Halchenko, John B. Cole, ... Adel Qalieh

Zenodo (2018-07-16) <https://doi.org/10.5281/zenodo.1313201>

47. pysam-developers/pysam

pysam-developers

GitHub <https://github.com/pysam-developers/pysam>

48. **The Sequence Alignment/Map format and SAMtools.**

Heng Li, Bob Handsaker, Alec Wysoker, Tim Fennell, Jue Ruan, Nils Homer, Gabor Marth, Goncalo Abecasis, Richard Durbin,

Bioinformatics (Oxford, England) (2009-06-08) <https://www.ncbi.nlm.nih.gov/pubmed/19505943>

49. **AMD, an Automated Motif Discovery Tool Using Stepwise Refinement of Gapped Consensuses**

Jiantao Shi, Wentao Yang, Mingjie Chen, Yanzhi Du, Ji Zhang, Kankan Wang

PLoS ONE (2011-09-12) <https://doi.org/10.1371/journal.pone.0024576>

50. **BioProspector: discovering conserved DNA motifs in upstream regulatory regions of co-expressed genes.**

X Liu, DL Brutlag, JS Liu

Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing (2001) <https://www.ncbi.nlm.nih.gov/pubmed/11262934>

51. **Deep and wide digging for binding motifs in ChIP-Seq data**

I. V. Kulakovskiy, V. A. Boeva, A. V. Favorov, V. J. Makeev

Bioinformatics (2010-10-15) <https://doi.org/10.1093/bioinformatics/btq488>

52. **GADEM: A Genetic Algorithm Guided Formation of Spaced Dyads Coupled with an EM Algorithm for Motif Discovery**

Leping Li

Journal of Computational Biology (2009-02) <https://doi.org/10.1089/cmb.2008.16tt>

53. **On the detection and refinement of transcription factor binding sites using ChIP-Seq data**

Ming Hu, Jindan Yu, Jeremy M. G. Taylor, Arul M. Chinnaiyan, Zhaohui S. Qin

Nucleic Acids Research (2010-01-07) <https://doi.org/10.1093/nar/gkp1180>

54. **cis-Site Seeker**<https://users.soe.ucsc.edu/~kent/improbizer/index.html>

55. **An algorithm for finding protein–DNA binding sites with applications to chromatin-immunoprecipitation microarray experiments**

X. Shirley Liu, Douglas L. Brutlag, Jun S. Liu

Nature Biotechnology (2002-07-08) <https://doi.org/10.1038/nbt717>

56. **Fitting a mixture model by expectation maximization to discover motifs in biopolymers.**

TL Bailey, C Elkan

Proceedings. International Conference on Intelligent Systems for Molecular Biology (1994) <https://www.ncbi.nlm.nih.gov/pubmed/7584402>

57. A Gibbs Sampling Method to Detect Overrepresented Motifs in the Upstream Regions of Coexpressed Genes

Gert Thijs, Kathleen Marchal, Magali Lescot, Stephane Rombauts, Bart De Moor, Pierre Rouzé, Yves Moreau

Journal of Computational Biology (2002-04) <https://doi.org/10.1089/10665270252935566>

58. A highly efficient and effective motif discovery method for ChIP-seq/ChIP-chip data using positional information

Xiaotu Ma, Ashwinikumar Kulkarni, Zhihua Zhang, Zhenyu Xuan, Robert Serfling, Michael Q. Zhang

Nucleic Acids Research (2011-01-06) <https://doi.org/10.1093/nar/gkr1135>

59. Trawler: de novo regulatory motif discovery pipeline for chromatin immunoprecipitation

Laurence Ettwiller, Benedict Paten, Mirana Ramialison, Ewan Birney, Joachim Wittbrodt

Nature Methods (2007-06-24) <https://doi.org/10.1038/nmeth1061>

60. Using Weeder, Pscan, and PscanChIP for the Discovery of Enriched Transcription Factor Binding Site Motifs in Nucleotide Sequences

Federico Zambelli, Graziano Pesole, Giulio Pavesi

Current Protocols in Bioinformatics (2014-09) <https://doi.org/10.1002/0471250953.bi0211s47>

61. P-value-based regulatory motif discovery using positional weight matrices

H. Hartmann, E. W. Guthohrlein, M. Siebert, S. Luehr, J. Soding

Genome Research (2012-09-18) <https://doi.org/10.1101/gr.139881.112>

62. Robust rank aggregation for gene list integration and meta-analysis

Raivo Kolde, Sven Laur, Priit Adler, Jaak Vilo

Bioinformatics (2012-01-12) <https://doi.org/10.1093/bioinformatics/btr709>

63. <https://www.jstor.org/stable/2346101>

64.

Szymon M Kielbasa, Didier Gonze, Hanspeter Herzel

BMC Bioinformatics (2005) <https://doi.org/10.1186/1471-2105-6-237>

65. Motif clustering with implications for transcription factor interactions

Jan Grau, Ivo Grosse, Stefan Posch, Jens Keilwagen

PeerJ (2015-08-13) <https://doi.org/10.7287/peerj.preprints.1302v1>

66. Design of shortest double-stranded DNA sequences covering all k-mers with applications to protein-binding microarrays and synthetic enhancers

Y. Orenstein, R. Shamir

Bioinformatics (2013-06-21) <https://doi.org/10.1093/bioinformatics/btt230>

67. **BEDTools: a flexible suite of utilities for comparing genomic features**

Aaron R. Quinlan, Ira M. Hall

Bioinformatics (2010-01-28) <https://doi.org/10.1093/bioinformatics/btq033>

68. **Snakemake—a scalable bioinformatics workflow engine**

J. Koster, S. Rahmann

Bioinformatics (2012-08-20) <https://doi.org/10.1093/bioinformatics/bts480>

69. **Kundajelab/Atac_Dnase_Pipelines: 0.3.0**

Jin Lee, Grey Christoforo, Grey Christoforo, CS Foo, Chris Probert, Anshul Kundaje, Nathan Boley, Kohpangwei, Daniel Kim, Mike Dacre

Zenodo (2016-09-27) <https://doi.org/10.5281/zenodo.156534>

70. **Fast gapped-read alignment with Bowtie 2**

Ben Langmead, Steven L Salzberg

Nature Methods (2012-03-04) <https://doi.org/10.1038/nmeth.1923>

71. **Model-based Analysis of ChIP-Seq (MACS)**

Yong Zhang, Tao Liu, Clifford A Meyer, Jérôme Eeckhoute, David S Johnson, Bradley E Bernstein, Chad Nussbaum, Richard M Myers, Myles Brown, Wei Li, X Shirley Liu

Genome Biology (2008) <https://doi.org/10.1186/gb-2008-9-9-r137>

72. **FROM BINDING MOTIFS IN CHIP-SEQ DATA TO IMPROVED MODELS OF TRANSCRIPTION FACTOR BINDING SITES**

IVAN KULAKOVSKIY, VICTOR LEVITSKY, DMITRY OSHCHEPKOV, LEONID BRYZGALOV, ILYA VORONTSOV, VSEVOLOD MAKEEV

Journal of Bioinformatics and Computational Biology (2013-02) <https://doi.org/10.1142/s0219720013400040>

73. **The Next Generation of Transcription Factor Binding Site Prediction**

Anthony Mathelier, Wyeth W. Wasserman

PLoS Computational Biology (2013-09-05) <https://doi.org/10.1371/journal.pcbi.1003214>

74. **Varying levels of complexity in transcription factor binding motifs**

Jens Keilwagen, Jan Grau

Nucleic Acids Research (2015-06-26) <https://doi.org/10.1093/nar/gkv577>

75. **InMoDe: tools for learning and visualizing intra-motif dependencies of DNA binding sites**

Ralf Eggeling, Ivo Grosse, Jan Grau

Bioinformatics (2016-12-28) <https://doi.org/10.1093/bioinformatics/btw689>

76. **Automated incorporation of pairwise dependency in transcription factor binding site prediction using dinucleotide weight tensors**

Saeed Omid, Mihaela Zavolan, Mikhail Pachkov, Jeremie Breda, Severin Berger, Erik van Nimwegen

PLOS Computational Biology (2017-07-28) <https://doi.org/10.1371/journal.pcbi.1005176>

77. Predicting effects of noncoding variants with deep learning-based sequence model

Jian Zhou, Olga G Troyanskaya

Nature Methods (2015-08-24) <https://doi.org/10.1038/nmeth.3547>

78. Sequential regulatory activity prediction across chromosomes with convolutional neural networks

David R. Kelley, Yakir A. Reshef, Maxwell Bileschi, David Belanger, Cory Y. McLean, Jasper Snoek

Genome Research (2018-03-27) <https://doi.org/10.1101/gr.227819.117>

79. Discovering epistatic feature interactions from neural network models of regulatory DNA sequences

Peyton Greenside, Tyler Shimko, Polly Fordyce, Anshul Kundaje

Bioinformatics (2018-09-01) <https://doi.org/10.1093/bioinformatics/bty575>