

DNA Shape Features Improve Transcription Factor Binding Site Predictions In Vivo

Anthony Mathelier,^{1,2,3} Beibei Xin,⁴ Tsu-Pei Chiu,⁴ Lin Yang,⁴ Remo Rohs,^{4,*} and Wyeth W. Wasserman^{1,*}

¹Centre for Molecular Medicine and Therapeutics at the Child and Family Research Institute, Department of Medical Genetics, University of British Columbia, 980 West 28th Avenue, Vancouver, BC V5Z 4H4, Canada

²Centre for Molecular Medicine Norway (NCMM), Nordic EMBL Partnership, University of Oslo and Oslo University Hospital, 0318 Oslo, Norway

³Department of Cancer Genetics, Institute for Cancer Research, Oslo University Hospital Radiumhospitalet, 0372 Oslo, Norway

⁴Molecular and Computational Biology Program, Departments of Biological Sciences, Chemistry, Physics, and Computer Science, University of Southern California, Los Angeles, CA 90089, USA

*Correspondence: rohs@usc.edu (R.R.), wyeth@cmmt.ubc.ca (W.W.W.)

<http://dx.doi.org/10.1016/j.cels.2016.07.001>

SUMMARY

Interactions of transcription factors (TFs) with DNA comprise a complex interplay between base-specific amino acid contacts and readout of DNA structure. Recent studies have highlighted the complementarity of DNA sequence and shape in modeling TF binding in vitro. Here, we have provided a comprehensive evaluation of in vivo datasets to assess the predictive power obtained by augmenting various DNA sequence-based models of TF binding sites (TFBSs) with DNA shape features (helix twist, minor groove width, propeller twist, and roll). Results from 400 human ChIP-seq datasets for 76 TFs show that combining DNA shape features with position-specific scoring matrix (PSSM) scores improves TFBS predictions. Improvement has also been observed using TF flexible models and a machine-learning approach using a binary encoding of nucleotides in lieu of PSSMs. Incorporating DNA shape information is most beneficial for E2F and MADS-domain TF families. Our findings indicate that incorporating DNA sequence and shape information benefits the modeling of TF binding under complex in vivo conditions.

INTRODUCTION

One of many mechanisms that control gene expression, transcriptional regulation involves transcription factors (TFs) as key proteins (Jacob and Monod, 1961; Ptashne and Gann, 1997). Most TFs are sequence-specific DNA binding proteins that recognize specific genome positions through a complex interplay between nucleotide-amino-acid contacts (base readout) and readout of DNA structure (shape readout) (Slattery et al., 2014). Deciphering how TFs identify and bind specific target sequences—the TF binding sites (TFBSs)—is a key challenge in understanding transcriptional gene regulation (Dror et al., 2016; Wasserman and Sandelin, 2004; Zambelli et al., 2013).

TFBSs are short and often degenerate sequence motifs. These characteristics make it computationally difficult to model and predict TFBSs at the genomic scale (Badis et al., 2009). Moving beyond initial consensus sequence methods, the classical computational model to describe TFBSs is the position-specific scoring matrix (PSSM), which uses an additive method to summarize frequencies of every nucleotide at each position of the TFBS (Stormo, 2013). These second-generation models, however, do not capture position interdependencies or variable spacing. Therefore, several experimental assays have been designed to unravel characteristics of TF-DNA interactions at the large scale. In vitro high-throughput (HT) binding assays, such as protein binding microarrays (PBMs) (Berger et al., 2006), HT SELEX (Jolma et al., 2010; Zhao et al., 2009), and SELEX-seq (Slattery et al., 2011), expose DNA sequences selected by TFs and reveal their binding preferences. Chromatin immunoprecipitation sequencing (ChIP-seq) represents the in vivo counterpart of these in vitro assays, allowing for the identification of DNA regions bound by a targeted TF at the genomic scale (Johnson et al., 2007).

Large-scale data derived from HT experiments highlight higher-order positional interaction features of TFBSs that cannot be captured by classical PSSMs, even though the methods based on these traditional models perform quite well (Weirauch et al., 2013). Recently, computational advances have used experimental assays to construct sophisticated models that capture a broad range of TFBS representations. For instance, PSSMs have been extended to dinucleotides to capture interrelationships within TFBSs (Siddharthan, 2010). Using PBM data, binding energy models include energy parameters to describe contributions of dinucleotides to binding affinity (Zhao et al., 2012). These models describe TF-DNA binding specificity well in cases in which PSSMs have performed insufficiently. Utilizing ChIP-seq data, we developed the TF flexible model (TFFM) framework to improve in vivo prediction of TFBSs (Mathelier and Wasserman, 2013). TFFMs capture interdependencies of successive nucleotides within TFBSs and the flexible length of TFBSs within a single hidden Markov model framework.

The above-mentioned third-generation methods enable TFBS prediction by representing sequence properties. A parallel approach utilizes the 3D DNA structure, or DNA shape, to capture, at least in part, the interdependencies between nucleotide

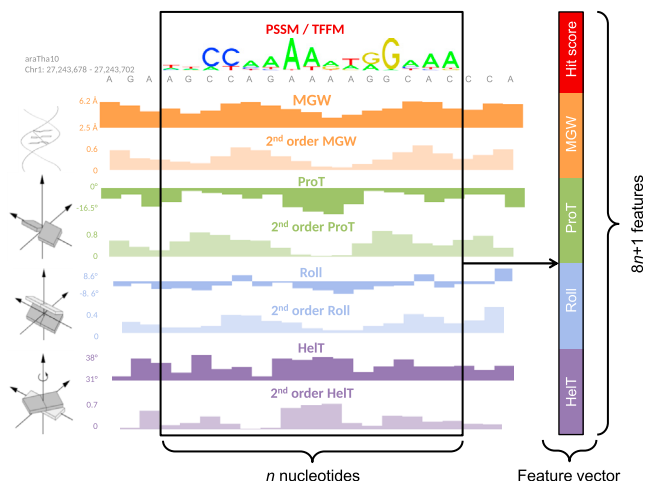


Figure 1. Feature Vectors of PSSM + DNA Shape and TFFM + DNA Shape Classifiers

Feature vectors combine sequence scores with respect to the TF binding profile (PSSM or TFFM), the normalized values of four DNA shape features (MGW, ProT, Roll, and HeIT), and their normalized product terms at adjacent positions as second-order shape features (Zhou et al., 2015). In 4-bits + DNA shape classifiers, the TF binding profile score is replaced by binary 4-bits encoding of the corresponding sequence (Zhou et al., 2015).

positions within TFBSs (Gordân et al., 2013; Tsai et al., 2015; Yang and Ramsey, 2015; Zhou et al., 2015). Large-scale DNA structural information can be computed by the DNashape method (Zhou et al., 2013), which computes four DNA shape features: helix twist (HeIT), minor groove width (MGW), propeller twist (ProT), and Roll. Recent studies have demonstrated the complementary role of DNA sequence and shape information in determining protein-DNA binding specificity in vitro (Joshi et al., 2007; Rohs et al., 2009; Slattery et al., 2011). For example, the binding specificity of Hox proteins was analyzed using SELEX-seq data to show the direct role of DNA shape features in protein-DNA readout (Abe et al., 2015). Using PBM and SELEX-seq data, we showed that complementing DNA sequence with shape information enhanced the prediction of TF binding affinities (Zhou et al., 2015). DNA shape information at regions flanking core binding sites was highly predictive of differential binding derived from BunDLE-seq assays (Levo et al., 2015).

While previous works have demonstrated that models combining DNA sequence and shape improve quantitative models of TF binding in vitro, we have addressed here three key questions:

- (1) Do more complex in vivo protein-DNA interactions exhibit similar properties?
- (2) When DNA shape properties are integrated with sequence-based TFBS prediction methods, do we observe an improvement in performance?
- (3) Do specific TF families benefit more than others from the integration of DNA shape features in TF binding models?

Here, we have capitalized on the availability of DNA shape information extracted from GBshape (Chiu et al., 2015), our genome browser database of DNA shape features computed from our DNashape prediction tool (Zhou et al., 2013), at

TF-bound regions derived from ChIP-seq experiments to address the three aforementioned questions.

RESULTS

Machine Learning Models Combining DNA Sequence and Shape Features

To assess the effects of including DNA structural information in predictions of TFBSs in ChIP-seq datasets, we developed a computational framework combining DNA sequence and shape information to model and predict TFBSs. The availability of numerous ChIP-seq regions enables the application of a discriminative supervised machine learning approach (Libbrecht and Noble, 2015). Specifically, a DNA sequence that is considered as a potential TFBS was represented by a (feature) vector that combined 1 to $4n$ features that encode sequence information and $8n$ features that capture DNA shape information, where n is DNA sequence length. We encoded DNA sequence information of the putative TFBS by using either the PSSM or TFFM score computed from the sequence, or a binary encoding using 4 bits per nucleotide (Zhou et al., 2015). DNA shape-related features are the predicted values of HeIT, MGW, ProT, and Roll at each position of the TFBS, extracted from GBshape (Chiu et al., 2015). The vector was further augmented with four second-order shape features that capture structural dependencies at adjacent nucleotide positions (Zhou et al., 2015) (Figure 1). Assuming that each ChIP-seq region contains a TFBS, we constructed a feature vector for the best hit per ChIP-seq peak and background region predicted by a TF binding profile (PSSM or TFFM) to train a classifier. To discriminate between TF bound (ChIP-seq) and unbound (background) regions, we used a gradient boosting classifier, which is an ensemble machine learning classifier that combines multiple weak learners to improve predictive power (Friedman et al., 2001). The gradient boosting classifier was based on decision trees that, given an input feature vector, output the probability of the feature vector to be associated with a ChIP-seq peak or a background region. This approach naturally handles heterogeneous features (e.g., DNA sequence and shape information), is robust to outliers, and is able to manage irrelevant input, such as noise from ChIP-seq experiments (Friedman et al., 2001).

Classifiers combining PSSM score, TFFM score, or 4-bits nucleotide encoding with DNA shape features are referred to as PSSM + DNA shape, TFFM + DNA shape, or 4-bits + DNA shape classifiers, respectively. Open-source Python software for generating and using these classifiers is provided at <https://github.com/amathelier/DNashapedTFBS>.

Incorporating DNA Shape Features Improves TFBS Prediction in Human In Vivo Datasets

We compiled a set of 400 uniformly processed human ENCODE ChIP-seq datasets for which a JASPAR TF binding profile (Mathelier et al., 2014) was available for the corresponding immunoprecipitated (ChIPed) TF (Data S1). These datasets, covering 76 TFs, were used to compare the predictive powers of three computational models that consider DNA sequence information alone with their DNA shape-augmented classifiers. The first two DNA sequence-based models are PSSMs and TFFMs, which are widely used to score TFBSs in ChIP-seq

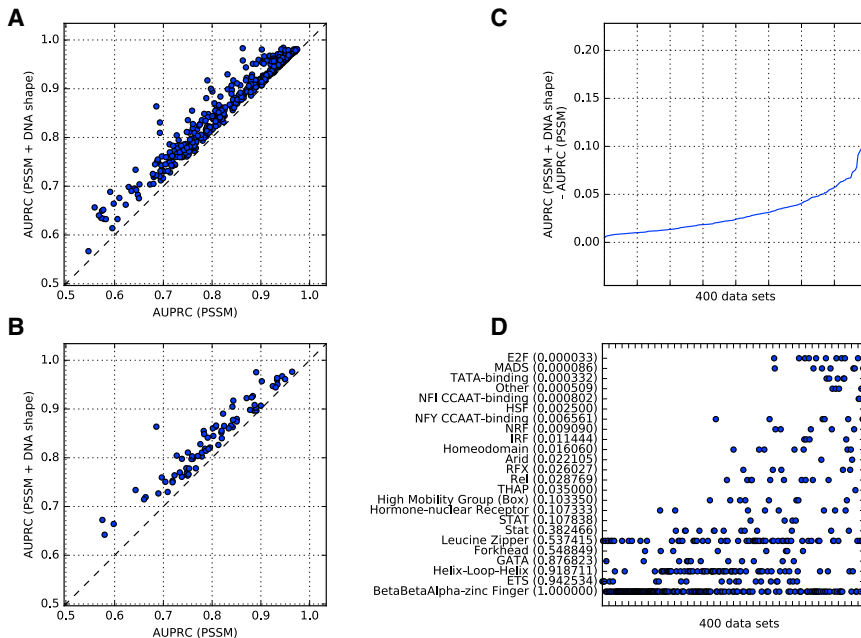


Figure 2. Effect of DNA Shape Features on TFBS Predictions in ChIP-Seq Data

(A) AUPRC values obtained for 400 ENCODE human ChIP-seq datasets using PSSM scores (x axis) or classifiers combining PSSM scores and DNA shape features (y axis). The dashed line represents equal AUPRC values obtained with both methods. (B) Median AUPRC values over all ChIP-seq datasets associated with each TF (one point per TF), obtained using PSSM scores (x axis) or PSSM + DNA shape classifiers (y axis). Dashed line represents equal AUPRC values obtained with both methods. (C) Predictive power improvement obtained when considering DNA shape features (y axis) as the difference between AUPRC values obtained with PSSM + DNA shape classifiers and PSSM scores. Larger difference corresponds to stronger improvement. Datasets (x axis) are ranked by increasing difference values. (D) For each TF family (y axis), an associated dataset is represented at the corresponding x-coordinate where the dataset appears in (B). Names of TF families are given on y-axis, with significant Mann-Whitney U test p values in parentheses (not corrected for multiple hypothesis testing). See also [Data S1](#), [S2](#), [S3](#), [S4](#), and [S7](#).

datasets. The third model, the 4-bits classifier, is a discriminative model that uses a binary encoding of DNA sequence information (Zhou et al., 2015).

Here, the predictive power of a model refers to its ability to discriminate ChIP-seq regions (defined as the 50-bp region surrounding each side of the ChIP-seq peak maximum) from matched background sequences. The 50-bp regions were selected because they are enriched for TFBSs (Worsley Hunt et al., 2014; Wilbanks and Facciotti, 2010). To avoid sequence composition biases, we selected each set of background sequences to match either the G+C (%GC) content or dinucleotide composition of the ChIP-seq regions. Unless otherwise indicated, background sequences matching the %GC content of ChIP-seq regions were used in the following results.

Predictive powers of PSSM scores and PSSM + DNA shape classifiers were assessed through 10-fold cross-validation (CV). We optimized the PSSMs derived from JASPAR TF binding profiles with the perceptron algorithm using the DiMO tool (Patel and Stormo, 2014) on the constructed foreground and background training sets (range: 495–83,123, median: 15,171, mean: 21,098, SD: 17,220 sequences). Parameters of PSSM + DNA shape classifiers were learned from the same training sets. Vectors used by the classifiers for a ChIP-seq region correspond to the combination of the best PSSM score in the region and the 8n DNA shape feature values computed for this hit. To assess predictive power, we varied the threshold for scores to compute the recall (sensitivity), specificity, and precision values. Areas under the precision and recall curve (AUPRC) and the receiver-operating characteristic curve (AUROC) were computed for each model on each ChIP-seq dataset to evaluate predictive power. Unless otherwise noted, we provide the AUPRC values and the p values for significance calculated by a Wilcoxon signed-rank test.

By comparing AUPRC values derived from the PSSM scores or PSSM + DNA shape classifiers, we found that shape-augmented classifiers performed better for all 400 ChIP-seq datasets ($p = 2.7 \times 10^{-67}$; Figure 2A). By considering the median AUPRC values per TF over all ChIP-seq datasets associated with the TF, we observed consistent improvement for all TFs when DNA shape features were incorporated ($p = 3.6 \times 10^{-14}$; Figure 2B). We computed the difference of discriminative power between the two models (Figure 2C) to assess the improvement obtained by using the PSSM + DNA shape classifiers.

Using the same analyses, we found that the predictive power of the TFFM + DNA shape classifiers is better than that of the TFFMs for 396/400 ChIP-seq datasets ($p = 4.4 \times 10^{-67}$; Data S2). Classifiers performed strictly better than TFFMs for all TFs when we considered the median AUPRC values per TF ($p = 3.6 \times 10^{-14}$; Data S2).

Finally, we compared the 4-bits and 4-bits + DNA shape classifiers, which were trained and tested on sequences of the highest-scoring hit per ChIP-seq region derived from the PSSMs. DNA shape-augmented classifiers performed consistently better than 4-bits classifiers for 365/400 ChIP-seq datasets ($p = 2.7 \times 10^{-57}$) and 70/76 TFs ($p = 1.3 \times 10^{-12}$) when considering the median AUPRC values (Data S2).

We confirmed the improvement in discriminative power of the models incorporating DNA shape features by considering background sequences matching the dinucleotide composition of ChIP-seq regions (Data S3) and TF-bound regions recurrently found in multiple ChIP-seq datasets for the same TF (Data S4).

The relative improvement obtained when incorporating DNA shape information varied depending on the baseline DNA sequence-based approach. Unsurprisingly, the 4-bits + DNA shape classifiers exhibited a smaller improvement over the 4-bits classifiers compared to the shape-based improvements

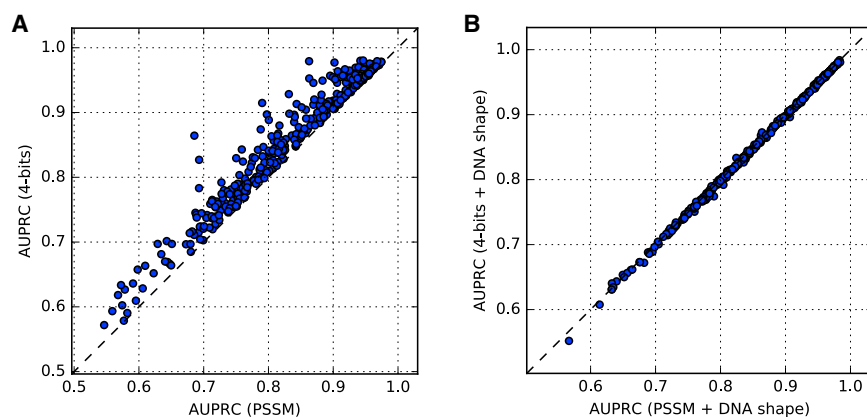


Figure 3. Predictive Power of PSSM and 4-Bit Approaches for TFBSs in ChIP-Seq Regions

(A and B) AUPRC values obtained for 400 human ENCODE ChIP-seq datasets, obtained by considering PSSM scores (x axis) and 4-bits classifiers (y axis) (A) or PSSM + DNA shape (x axis) and 4-bits + DNA shape (y axis) classifiers (B). See also [Data S5](#).

obtained with PSSMs and TFFMs. The higher baseline performance of the 4-bits method is consistent with the superiority of discriminative over generative models to distinguish bound from unbound regions in ChIP-seq (Libbrecht and Noble, 2015) (Figure 3A; Data S5). Nonetheless, PSSM + DNA shape classifiers performed consistently better than 4-bits + DNA shape classifiers for 344/400 datasets ($p = 7.7 \times 10^{-43}$; Figure 3B) and 64/76 TFs ($p = 1.0 \times 10^{-8}$; Data S5).

Although 4-bits classifiers outperformed PSSM scores, the higher discriminative power of PSSM + DNA shape compared to 4-bits + DNA shape classifiers reinforces the capacity of DNA shape features to improve TFBS predictions in ChIP-seq datasets. Importantly, the combination of sequence information (captured by PSSMs, TFFMs, or 4-bits classifiers) with DNA shape properties performed better than generative (PSSM and TFFM) and discriminative (4-bits classifier) approaches modeling DNA sequence, indicating that DNA shape provides additional information.

Although the utility of DNA shape to predict TFBSs was reported before (Abe et al., 2015; Yang and Ramsey, 2015; Yang et al., 2014; Zhou et al., 2015), we provide evidence, from an extensive collection of 400 human in vivo datasets for 76 TFs, that this observation is generalizable and relevant to noisy environments and data (Fan and Struhl, 2009; Worsley Hunt et al., 2014; Worsley Hunt and Wasserman, 2014; Jain et al., 2015; Park et al., 2013; Teytelman et al., 2013).

DNA Shape at Genomic Flanking Regions Improves TFBS Predictions In Vivo

Sequences immediately flanking TFBSs have been previously shown to contribute to TF binding specificity (Gordân et al., 2013), which is determined, in part, by DNA shape outside the core binding sites (Afek et al., 2014; Barozzi et al., 2014; Dror et al., 2015). We extended our DNA shape-augmented models to consider eight DNA shape features at 15-bp-long regions 5' and 3' of the TFBSs, as in Barozzi et al. (2014).

Augmenting DNA shape-based classifiers with additional DNA shape information from flanking sequences has improved the discriminatory power of classifiers trained using 10-fold CV for 378 (~94%), 373 (~93%), and 375 (~94%) datasets compared to PSSM + DNA shape, TFFM + DNA shape, and 4-bits + DNA shape classifiers, respectively (Figure 4; Data S6). Our findings agree with results from in vitro studies of the role of flanking regions in TF-DNA binding (Dror et al., 2016; Gordân et al., 2013; Levo et al., 2015).

E2F and MADS-Domain TF Families Benefit Most from DNA Shape Information

Next, motivated by the observation that DNA structural information improves the prediction of TFBSs for some ChIP-seq datasets more than others (Figure 2C; Data S2 and S3), we investigated whether predictions for certain TF families with similar DNA-binding domains specifically benefit from incorporating DNA shape information.

By using JASPAR (Mathelier et al., 2014), we extracted TF family information of DNA binding domains associated with the




Figure 4. Predictive Power of DNA Shape Features at TFBS Flanking Regions

(A) AUPRC values obtained for 400 human ENCODE ChIP-seq datasets when using classifiers combining PSSM scores and DNA shape features at core TFBSs (x axis), or classifiers combining PSSM scores and DNA shape features at both core TFBSs and surrounding 15 bp on each side (y axis). Dashed line represents equal AUPRC values for both methods.

(B) AUPRC value differences (y axis) between flank-augmented classifiers and PSSM + DNA shape classifiers (x axis). Datasets (x axis) are ranked by increasing difference values. See also [Data S6](#).




Figure 5. Use of a Single DNA Shape Feature Category for E2F and MADS-Box TFBS Recognition in ChIP-Seq

(A and B) AUPRC values (y axis) for E2F (A) and MADS-domain (B) TF datasets (x axis), obtained by using all four first-order DNA shape features or a single feature category along with sequence features in the PSSM + DNA shape classifiers.

To confirm the results obtained when incorporating DNA shape features for MADS-domain TFs, we considered an independent set of seven *Arabidopsis thaliana* TFs (Heyndrickx et al., 2014), for which we had JASPAR TF binding profiles (Data S1). Similar to human MADS-domain TF results (Data S7), we observed improved discriminative power when DNA structural information was considered for plant MADS-domain TFs with PSSM + DNA shape and TFFM + DNA shape models (Data S7). Only the two smallest datasets (94 sequences for FLC and 54 for SVP in training sets) showed decreased discriminative power with DNA shape-augmented classifiers. For one of the three approaches, the 4-bits + DNA shape classifiers, we found no improvement in predictive power for plant MADS-domain TF datasets compared to the 4-bits classifiers (Data S7).

400 human ChIP-seq experiments. In aggregate, we analyzed profiles derived from DNA binding domains associated with 24 TF families, which were associated with TFs in the JASPAR database using a classification scheme (Fulton et al., 2009) (Data S1). By comparing the predictive powers of DNA shape-augmented classifiers to those of DNA sequence-based approaches, we assessed the enrichment of a TF family for larger AUPRC or AUROC difference values using the one-sided Mann-Whitney U test. Predictive power comparisons were performed considering %GC- and dinucleotide-matched background sets.

Depending on the DNA sequence-based approach (PSSM, TFFM, or 4-bits), background type (%GC- or dinucleotide-matched), and assessment method (AUPRC or AUROC), we observed enrichment ($p < 4.17 \times 10^{-4}$, with Bonferroni correction for desired $\alpha < 0.01$) for different TF families (Figure 2D; Data S1 and S2). Taken together, the results across all three DNA sequence-based approaches suggest that the E2F and MADS-domain TF datasets benefited the most from inclusion of DNA shape information (Figure 2D; Data S1 and S2). It is noteworthy that results for the E2F and MADS-domain TF datasets were not derived from a single TF, but were consistent over several TFs from the same family. Namely, the ten E2F-associated ChIP-seq datasets were derived from E2F1 (three datasets), E2F4 (four datasets), and E2F6 (three datasets), and seven MADS-domain-associated ChIP-seq datasets were derived from MEF2A (two datasets), MEF2C (one dataset), and SRF (four datasets).

Taken together, we observed that, among all studied protein families, TFs from the E2F and MADS-domain families most benefited from the inclusion of DNA shape information in the classifiers when compared to the three DNA-sequence-based models.

Structural Analyses of E2F and MADS-Domain TF Binding Specificities

DNA Shape Readout Contributes to E2F-DNA Binding

Given our observed improvement in predictive power based on DNA shape information, we next characterized the specific DNA shape features that contributed the most to the improvement in predictive accuracy for the family of E2F TFs. We extracted the importance of each feature learned by the DNA shape-augmented classifiers by combining DNA sequence and shape information for the E2F TFs. To consider the same DNA sequence-based model per TF for all associated ChIP-seq datasets, we selected PSSMs derived from the corresponding JASPAR TF binding profiles. To simplify interpretation, we considered classifiers based on sequence and first-order DNA shape features (Zhou et al., 2015). Figures S1 and S2 plot the average feature importance measures obtained over the 10-fold CV training for all ChIP-seq datasets associated with the TFs.

Although the PSSM score was consistently the most important feature (Figure S1), several DNA shape features at different positions were important for TFBS predictions. A commonality

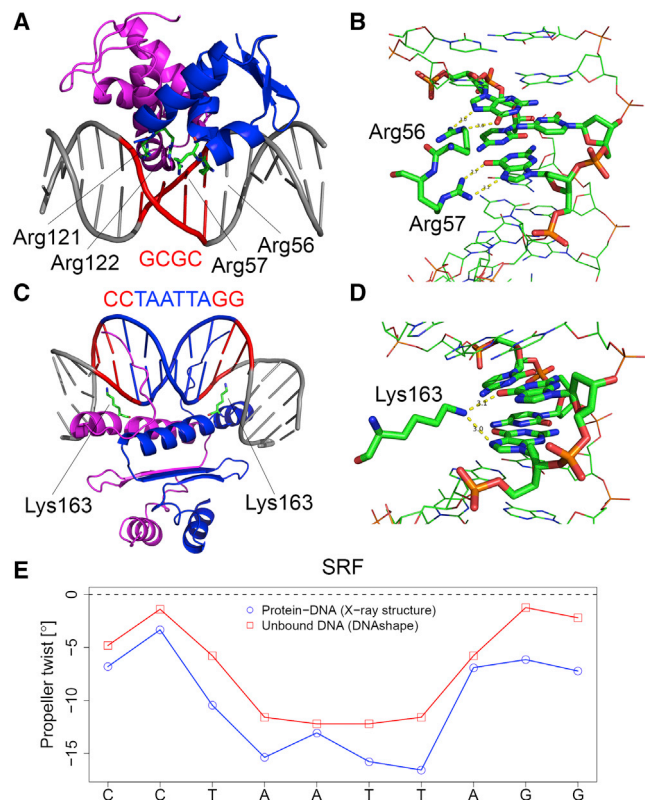


Figure 6. Structural Analysis of E2F and MADS-Domain TFs in Complex with DNA

(A) Co-crystal structure (PDB: 1CF7) of E2F4 (blue) and DP2 (magenta) forming a heterodimer that binds to core motif GCGC (red).
 (B) Detailed view of hydrogen bonds between arginines and guanines in major groove.
 (C) Co-crystal structure (PDB: 1SRS) of MADS-domain SRF homodimer in complex with core motif CCTAATTAGG.
 (D) Detailed view of hydrogen bonds between lysine and guanines in major groove.
 (E) ProT in bound (blue; calculated from X-ray structure) and unbound (red; predicted by DNASHape) target site.

among the three E2Fs was the contribution of ProT (and, to lesser extents, HelT and MGW) at proximal flanks of the TFBSs (see first and last positions of heat maps in Figure S2). Nucleotides immediately flanking the E2F TFBSs contributed to the DNA binding specificity of this TF family.

Comparing the predictive powers of classifiers combining sequence with either a single or four DNA shape features, we confirmed the importance of ProT for improving predictive power for the ten ChIP-seq datasets associated with the three E2Fs (Figures 5A and S3A).

Analysis of the co-crystal structure of the E2F4 TF in complex with DNA and its cofactor DP2 (PDB ID 1CF7) (Zheng et al., 1999) revealed that the RXXYD motifs of the E2F4 and DP2 heterodimer form a compact structural assembly that contacts the major groove (Figures 6A and 6B). Guanidinium groups of each of the four arginine residues engage in base readout through bidentate hydrogen bonds with guanine bases of the core binding site. This intricate system of eight hydrogen bonds enables readout of structural features and recognition of functional

groups of the guanine bases. The angular orientation of the arginine side chains stabilized by other amino acids selects for rotational parameters of the contacted G/C bp through hydrogen bond geometries. This observation is reflected by the importance given to rotational parameters, such as ProT, in our models. Similar results were reported in a recent study of diverse protein families (Dror et al., 2015).

MADS-Domain TFs Recognize Position-Specific DNA Shape

As described above for E2Fs, we individually considered DNA shape features in the DNA shape-augmented classifiers for the MADS-domain TF ChIP-seq datasets. Figure 5B highlights the importance of ProT for improving the discriminative power of models associated with human and plant MADS-domain TFs. The FLC and SVP datasets, which had small numbers of training sequences, were the only ones for which ProT was not the most important shape feature. Inclusion of a single DNA shape category ensured that DNA shape features did not compensate for each other when considering all four shape features due to dependencies among different features. Previous work showed the important role of MGW at the A-tract of the MADS-box for DNA-binding (Muiño et al., 2014). Our models captured this importance, although ProT remained the most important shape feature.

We extracted feature importance measures at each position within the TFBSs learned by the PSSM + DNA shape classifiers for human and plant MADS-domain TFs (Figures S4, S4, and S6). The most important DNA shape features for discriminating ChIP-seq bound sites from background genomic regions were ProT and Roll at specific positions within the MADS-box TFBSs (in agreement with Figure 5B). This observation was consistent across all of the human MADS-domain TFs, whereas the signal was more diffuse for plant MADS-domain TFs. DNA shape-augmented classifiers associated with human MADS-domain TF ChIP-seq datasets obtained the strongest discriminative improvements over sequence-based models (Figure 7A).

As an example, we plotted feature importance measures and sequence logos of the TF binding profiles associated with SRF and MEF2C (Figures 7B and 7C). These two TFs were associated with datasets that showed the strongest improvements in discriminative power when incorporating DNA shape features for the classification of bound vs. unbound sites (Figure 7A). ProT features seemed to contribute to the binding of the core CARG-box (CCW6GG) (red squares in Figures 7B, 7C, and S7), whereas Roll features were highlighted at the edges of the MADS-box core motif (blue squares in Figures 7B, 7C, and S7).

Structural analyses of various complexes of MADS-domain TFs with their DNA target sites suggested that protein residues recognize specific DNA conformations. Comparison of the feature importance measures for different structural features of the DNA binding sites of human MADS-domain proteins indicated a contribution of ProT to binding specificity (Figures 6, 7, S3B, and S7). Human SRF ChIP-seq datasets were consistently associated with strongest improvement in discriminative power when considering the DNA shape-augmented classifiers to predict TFBSs in TF-bound regions (Figure 7A). Hence, we analyzed the co-crystal structures available for complexes of human SRF and MEF2.

The SRF homodimer uses lysine residues to form base-specific hydrogen bonds with C/G bp in the CARG motif (Figures 6C and

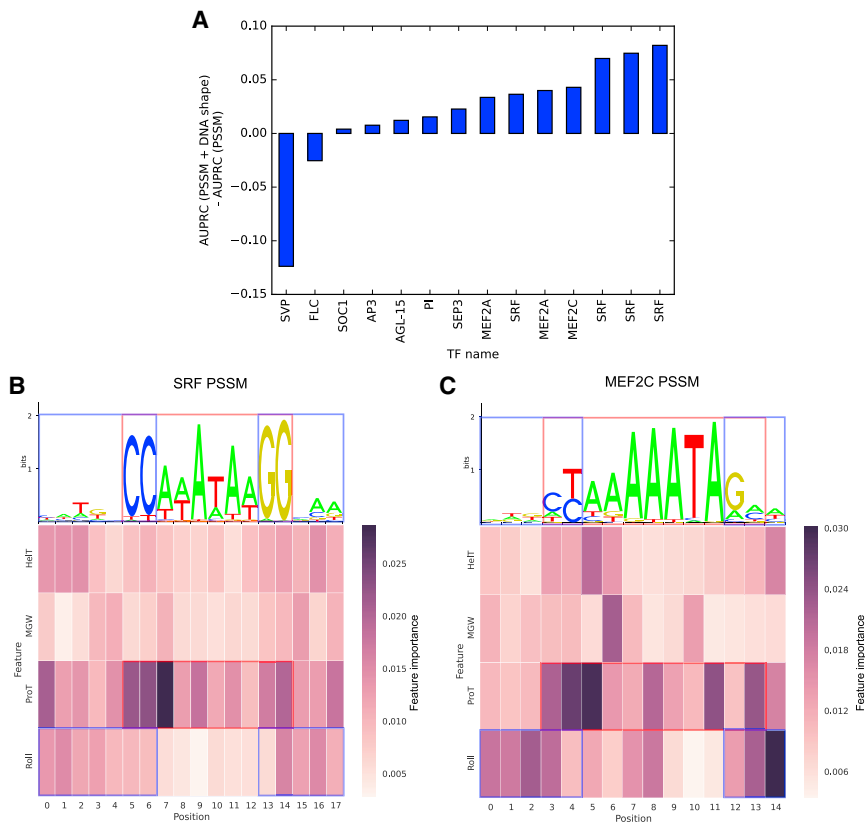


Figure 7. Feature Importance Measures for MADS-Box Recognition in ChIP-Seq Datasets

(A) AUPRC improvements (y axis) for human and plant MADS-domain TF ChIP-seq datasets (x axis provides TF names) when using PSSM + DNA shape features versus PSSM scores.

(B and C) Weblogos derived from JASPAR TF binding profile associated with SRF (MA0083.2) (B) and MEF2C (MA0497.1) (C) TFs are provided in top panels. Heat maps illustrating average feature importance values (y axis) at each position (x axis) of TFBSs in the classifiers trained for 10-fold CV analysis of ChIP-seq datasets are provided in bottom panels. Only feature importance measures associated with first-order DNA shape features are considered. Color scale for heat map is given on the right of the heat map. Red boxes highlight core MADS-box motif (CCW₆GG). Blue boxes highlight edges of motifs.

See also Figures S1, S2, and S4–S7.

6D). DNA bends around the protein, and intrinsic structural features are likely responsible for this deformation. As our models indicate a contribution of ProT, we compared ProT observed in the co-crystal structure of the SRF-DNA complex (PDB: 1SR5) (Pellegrini et al., 1995) with the prediction of ProT for the unbound motif using our DNashape method (Zhou et al., 2013). The ProT pattern in the unbound target site resembled the pattern in the protein-DNA complex (Figure 6E).

This observation suggests that ProT is intrinsic to the binding site and likely selected by the TF. A less negative ProT within the two adjacent G/C bp at each flank of the motif optimizes the geometry of hydrogen bonds between the lysine side chains and bases (Figure 6D). This preference is coupled to the sequence and selection of functional groups of the bases, whereas structural features enhance the energetics of these specific contacts.

DISCUSSION

Here, we used a machine learning classifier-based approach to demonstrate that combining DNA sequence information with DNA shape features improves discrimination between TF-bound sites in vivo and background genomic regions. These in vivo analyses complement our previous in vitro studies showing that inclusion of DNA shape properties can improve the accuracy of TF binding site prediction (Abe et al., 2015; Yang et al., 2014; Zhou et al., 2015).

A possible limitation of our approach is that we only considered DNA shape features at the best TFBS per ChIP-seq region,

derived from PSSM or TFFM scores. Although the site with the highest score represents the best candidate in a ChIP-seq region, another site harboring a lower score could potentially represent a more appropriate DNA shape readout (Zentner et al., 2015).

In our analyses, we used three baseline approaches (PSSM, TFFM, and 4-bits)

representing two types of models (generative and discriminative), two background types (%GC- and dinucleotide-matched), and two assessment measures (AUPRC and AUROC). Direct comparisons between the models stressed the higher predictive power of discriminative models using ChIP-seq datasets, in agreement with the literature (Libbrecht and Noble, 2015).

Our computational analyses of 400 human ChIP-seq datasets revealed that when DNA shape features were incorporated in the models, the E2F and MADS-domain TF families showed the largest improvement in TFBS prediction accuracies. TF families with the strongest predictive power improvements varied depending on the baseline model, background type, and assessment measure considered. These results highlight the importance of considering multiple background types and assessment measures when comparing models.

Whereas most bioinformatics tools rely on PSSMs for TFBS predictions, our findings imply that the field should consider more sophisticated modeling methods. For TFs for which only a small number of experimentally derived TF-bound regions are available, traditional PSSMs represent a reasonable alternative, as exemplified by the FLC and SVP plant TFs in our study. We envision that future tools relying on TFBS predictions will incorporate the most appropriate model for each TF or TF family.

With the increasing trove of whole-genome sequencing data, the identification of variants altering gene regulation through disruption of TF-DNA interactions has become an important challenge. Recent approaches have focused on allelic differences in PSSM scores to predict the functional impact of

variants within TFBSs (Chen et al., 2016; Mathelier et al., 2015). Our study confirms that some TFs critically rely on DNA shape readout for TFBS recognition. Future work should assess how the incorporation of DNA structural properties could help to predict the impact of variants disrupting TFBSs.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- METHOD DETAILS
 - ChIP-Seq Datasets and TF Binding Profiles
 - ChIP-Seq Peak Sequences
 - Recurrent ChIP-Seq Peak Regions
 - 10-fold CV Datasets
 - Discriminative Power Assessment
 - DNA Shape Features
 - Gradient Boosting Classifiers
 - Scanning DNA Sequences with TFFMs and PSSMs
 - TF Family Annotation
 - Feature Importance Measures
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Wilcoxon Signed-Rank Tests
 - Mann-Whitney U Tests
- DATA AND SOFTWARE AVAILABILITY
 - Software

SUPPLEMENTAL INFORMATION

Supplemental Information includes seven figures and seven data files and can be found with this article online at <http://dx.doi.org/10.1016/j.cels.2016.07.001>.

AUTHOR CONTRIBUTIONS

A.M., W.W.W., and R.R. conceived and designed the project. A.M. implemented and performed the experiments. A.M., W.W.W., and R.R. analyzed and interpreted the results. A.M. implemented the DNAsHapedTFBS tool. B.X. performed experiments and analyzed data related to the 4-bits-based classifiers, with contributions from T.-P.C. T.-P.C. also generated tools for DNA shape analysis. L.Y. analyzed and interpreted structural data. A.M., W.W.W., and R.R. wrote the manuscript.

ACKNOWLEDGMENTS

We thank Miroslav Hatas for systems support, Dora Pak for management support, François Parcy, Oriol Fornés Crespo, and Chih-Yu Chen for comments on the manuscript, and all members of the W.W.W. and R.R. labs for insightful discussions. A.M. and W.W.W. were supported by the Genome Canada Large scale Applied Research Grant (174CDE). Funding was provided by the Child and Family Research Institute, Vancouver, and the British Columbia Children's Hospital Foundation (to A.M. and W.W.W.). A.M. was also supported by funding from the Norwegian Research Council, Helse Sør-Øst, and the University of Oslo through the Centre for Molecular Medicine Norway (NCMM), which is a part of the Nordic European Molecular Biology Laboratory partnership for Molecular Medicine. R.R. was supported by the NIH (grants R01GM106056 and U01GM103804) and an Alfred P. Sloan Research Fellowship. Open access charges were defrayed in part through the National Science Foundation (grant MCB-1413539 to R.R.).

Received: September 12, 2015

Revised: March 4, 2016

Accepted: June 30, 2016

Published: August 18, 2016

REFERENCES

- Abe, N., Dror, I., Yang, L., Slattery, M., Zhou, T., Bussemaker, H.J., Rohs, R., and Mann, R.S. (2015). Deconvolving the recognition of DNA shape from sequence. *Cell* 161, 307–318.
- Afek, A., Schipper, J.L., Horton, J., Gordán, R., and Lukatsky, D.B. (2014). Protein-DNA binding in the absence of specific base-pair recognition. *Proc. Natl. Acad. Sci. USA* 111, 17140–17145.
- Badis, G., Berger, M.F., Philippakis, A.A., Talukder, S., Gehrke, A.R., Jaeger, S.A., Chan, E.T., Metzler, G., Vedenko, A., Chen, X., et al. (2009). Diversity and complexity in DNA recognition by transcription factors. *Science* 324, 1720–1723.
- Barozzi, I., Simonatto, M., Bonifacio, S., Yang, L., Rohs, R., Ghisletti, S., and Natoli, G. (2014). Coregulation of transcription factor binding and nucleosome occupancy through DNA features of mammalian enhancers. *Mol. Cell* 54, 844–857.
- Berger, M.F., Philippakis, A.A., Qureshi, A.M., He, F.S., Estep, P.W., 3rd, and Bulyk, M.L. (2006). Compact, universal DNA microarrays to comprehensively determine transcription-factor binding site specificities. *Nat. Biotechnol.* 24, 1429–1435.
- Chen, J., Rozowsky, J., Galeev, T.R., Harmanci, A., Kitchen, R., Bedford, J., Abyzov, A., Kong, Y., Regan, L., and Gerstein, M. (2016). A uniform survey of allele-specific binding and expression over 1000-Genomes-Project individuals. *Nat. Commun.* 7, 11101.
- Chiu, T.-P., Yang, L., Zhou, T., Main, B.J., Parker, S.C.J., Nuzhdin, S.V., Tullius, T.D., and Rohs, R. (2015). GBshape: a genome browser database for DNA shape annotations. *Nucleic Acids Res.* 43, D103–D109.
- Chiu, T.-P., Comoglio, F., Zhou, T., Yang, L., Paro, R., and Rohs, R. (2016). DNAsHapeR: an R/Bioconductor package for DNA shape prediction and feature encoding. *Bioinformatics* 32, 1211–1213.
- Cock, P.J., Antao, T., Chang, J.T., Chapman, B.A., Cox, C.J., Dalke, A., Friedberg, I., Hamelryck, T., Kauff, F., Wilczynski, B., and de Hoon, M.J. (2009). Biopython: freely available Python tools for computational molecular biology and bioinformatics. *Bioinformatics* 25, 1422–1423.
- Cunningham, F., Amode, M.R., Barrell, D., Beal, K., Billis, K., Brent, S., Carvalho-Silva, D., Clapham, P., Coates, G., Fitzgerald, S., et al. (2015). Ensembl 2015. *Nucleic Acids Res.* 43, D662–D669.
- Dror, I., Golan, T., Levy, C., Rohs, R., and Mandel-Gutfreund, Y. (2015). A widespread role of the motif environment in transcription factor binding across diverse protein families. *Genome Res.* 25, 1268–1280.
- Dror, I., Rohs, R., and Mandel-Gutfreund, Y. (2016). How motif environment influences transcription factor search dynamics: Finding a needle in a haystack. *BioEssays* 38, 605–612.
- Dunham, I., Kundaje, A., Aldred, S.F., Collins, P.J., Davis, C.A., Doyle, F., Epstein, C.B., Fritze, S., Harrow, J., Kaul, R., et al.; ENCODE Project Consortium (2012). An integrated encyclopedia of DNA elements in the human genome. *Nature* 489, 57–74.
- Fan, X., and Struhl, K. (2009). Where does mediator bind in vivo? *PLoS ONE* 4, e5029.
- Friedman, J., Hastie, T., and Tibshirani, R. (2001). The elements of statistical learning. Springer series in statistics (Springer).
- Fulton, D.L., Sundararajan, S., Badis, G., Hughes, T.R., Wasserman, W.W., Roach, J.C., and Sladek, R. (2009). TFCat: the curated catalog of mouse and human transcription factors. *Genome Biol.* 10, R29.
- Gordán, R., Shen, N., Dror, I., Zhou, T., Horton, J., Rohs, R., and Bulyk, M.L. (2013). Genomic regions flanking E-box binding sites influence DNA binding specificity of bHLH transcription factors through DNA shape. *Cell Rep.* 3, 1093–1104.

- Heyndrickx, K.S., Van de Velde, J., Wang, C., Weigel, D., and Vandepoele, K. (2014). A functional and evolutionary perspective on transcription factor binding in *Arabidopsis thaliana*. *Plant Cell* 26, 3894–3910.
- Jacob, F., and Monod, J. (1961). Genetic regulatory mechanisms in the synthesis of proteins. *J. Mol. Biol.* 3, 318–356.
- Jain, D., Baldi, S., Zabel, A., Straub, T., and Becker, P.B. (2015). Active promoters give rise to false positive ‘Phantom Peaks’ in ChIP-seq experiments. *Nucleic Acids Res.* 43, 6959–6968.
- Johnson, D.S., Mortazavi, A., Myers, R.M., and Wold, B. (2007). Genome-wide mapping of *in vivo* protein-DNA interactions. *Science* 316, 1497–1502.
- Jolma, A., Kivioja, T., Toivonen, J., Cheng, L., Wei, G., Enge, M., Taipale, M., Vaquerizas, J.M., Yan, J., Sillanpää, M.J., et al. (2010). Multiplexed massively parallel SELEX for characterization of human transcription factor binding specificities. *Genome Res.* 20, 861–873.
- Joshi, R., Passner, J.M., Rohs, R., Jain, R., Sosinsky, A., Crickmore, M.A., Jacob, V., Aggarwal, A.K., Honig, B., and Mann, R.S. (2007). Functional specificity of a Hox protein mediated by the recognition of minor groove structure. *Cell* 131, 530–543.
- Levo, M., Zalckvar, E., Sharon, E., Dantas Machado, A.C., Kalma, Y., Lotam-Pompan, M., Weinberger, A., Yakhini, Z., Rohs, R., and Segal, E. (2015). Unraveling determinants of transcription factor binding outside the core binding site. *Genome Res.* 25, 1018–1029.
- Libbrecht, M.W., and Noble, W.S. (2015). Machine learning applications in genetics and genomics. *Nat. Rev. Genet.* 16, 321–332.
- Mathelier, A., and Wasserman, W.W. (2013). The next generation of transcription factor binding site prediction. *PLoS Comput. Biol.* 9, e1003214.
- Mathelier, A., Zhao, X., Zhang, A.W., Parcy, F., Worsley-Hunt, R., Arenillas, D.J., Buchman, S., Chen, C.Y., Chou, A., Ienasescu, H., et al. (2014). JASPAR 2014: an extensively expanded and updated open-access database of transcription factor binding profiles. *Nucleic Acids Res.* 42, D142–D147.
- Mathelier, A., Shi, W., and Wasserman, W.W. (2015). Identification of altered cis-regulatory elements in human disease. *Trends Genet.* 31, 67–76.
- Muñoz, J.M., Smaczniak, C., Angenent, G.C., Kaufmann, K., and van Dijk, A.D.J. (2014). Structural determinants of DNA recognition by plant MADS-domain transcription factors. *Nucleic Acids Res.* 42, 2138–2146.
- Park, D., Lee, Y., Bhupindersingh, G., and Iyer, V.R. (2013). Widespread misinterpretable ChIP-seq bias in yeast. *PLoS ONE* 8, e83506.
- Patel, R.Y., and Stormo, G.D. (2014). Discriminative motif optimization based on perceptron training. *Bioinformatics* 30, 941–948.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Pellegrini, L., Tan, S., and Richmond, T.J. (1995). Structure of serum response factor core bound to DNA. *Nature* 376, 490–498.
- Perez, F., and Granger, B.E. (2007). IPython: a system for interactive scientific computing. *Comput. Sci. Eng.* 9, 21–29.
- Pohl, A., and Beato, M. (2014). bwtool: a tool for bigWig files. *Bioinformatics* 30, 1618–1619.
- Worsley Hunt, R., and Wasserman, W.W. (2014). Non-targeted transcription factors motifs are a systemic component of ChIP-seq datasets. *Genome Biol.* 15, 412.
- Worsley Hunt, R., Mathelier, A., Del Peso, L., and Wasserman, W.W. (2014). Improving analysis of transcription factor binding sites within ChIP-seq data based on topological motif enrichment. *BMC Genomics* 15, 472.
- Ptashne, M., and Gann, A. (1997). Transcriptional activation by recruitment. *Nature* 386, 569–577.
- Quinlan, A.R., and Hall, I.M. (2010). BEDTools: a flexible suite of utilities for comparing genomic features. *Bioinformatics* 26, 841–842.
- Rohs, R., West, S.M., Sosinsky, A., Liu, P., Mann, R.S., and Honig, B. (2009). The role of DNA shape in protein-DNA recognition. *Nature* 461, 1248–1253.
- Siddharthan, R. (2010). Dinucleotide weight matrices for predicting transcription factor binding sites: generalizing the position weight matrix. *PLoS ONE* 5, e9722.
- Slattery, M., Riley, T., Liu, P., Abe, N., Gomez-Alcala, P., Dror, I., Zhou, T., Rohs, R., Honig, B., Bussemaker, H.J., and Mann, R.S. (2011). Cofactor binding evokes latent differences in DNA binding specificity between Hox proteins. *Cell* 147, 1270–1282.
- Slattery, M., Zhou, T., Yang, L., Dantas Machado, A.C., Gordân, R., and Rohs, R. (2014). Absence of a simple code: how transcription factors read the genome. *Trends Biochem. Sci.* 39, 381–399.
- Stormo, G.D. (2013). Modeling the specificity of protein-DNA interactions. *Quant. Biol.* 1, 115–130.
- Teytelman, L., Thurtle, D.M., Rine, J., and van Oudenaarden, A. (2013). Highly expressed loci are vulnerable to misleading ChIP localization of multiple unrelated proteins. *Proc. Natl. Acad. Sci. USA* 110, 18602–18607.
- Tsai, Z.T.-Y., Shiu, S.-H., and Tsai, H.-K. (2015). Contribution of sequence motif, chromatin state, and dna structure features to predictive models of transcription factor binding in yeast. *PLoS Comput. Biol.* 11, e1004418.
- Wasserman, W.W., and Sandelin, A. (2004). Applied bioinformatics for the identification of regulatory elements. *Nat. Rev. Genet.* 5, 276–287.
- Weirauch, M.T., Cote, A., Norel, R., Annala, M., Zhao, Y., Riley, T.R., Saez-Rodriguez, J., Cokelaer, T., Vedenko, A., Talukder, S., et al.; DREAM5 Consortium (2013). Evaluation of methods for modeling transcription factor sequence specificity. *Nat. Biotechnol.* 31, 126–134.
- Wilbanks, E.G., and Facciotti, M.T. (2010). Evaluation of algorithm performance in ChIP-seq peak detection. *PLoS ONE* 5, e11471.
- Yang, J., and Ramsey, S.A. (2015). A DNA shape-based regulatory score improves position-weight matrix-based recognition of transcription factor binding sites. *Bioinformatics* 31, 3445–3450.
- Yang, L., Zhou, T., Dror, I., Mathelier, A., Wasserman, W.W., Gordân, R., and Rohs, R. (2014). TFBSshape: a motif database for DNA shape features of transcription factor binding sites. *Nucleic Acids Res.* 42, D148–D155.
- Zambelli, F., Pesole, G., and Pavesi, G. (2013). Motif discovery and transcription factor binding sites before and after the next-generation sequencing era. *Brief. Bioinform.* 14, 225–237.
- Zentner, G.E., Kasinathan, S., Xin, B., Rohs, R., and Henikoff, S. (2015). ChEC-seq kinetics discriminates transcription factor binding sites by DNA sequence and shape *in vivo*. *Nat. Commun.* 6, 8733.
- Zhao, Y., Granas, D., and Stormo, G.D. (2009). Inferring binding energies from selected binding sites. *PLoS Comput. Biol.* 5, e1000590.
- Zhao, Y., Ruan, S., Pandey, M., and Stormo, G.D. (2012). Improved models for transcription factor binding site identification using nonindependent interactions. *Genetics* 191, 781–790.
- Zheng, N., Fraenkel, E., Pabo, C.O., and Pavletich, N.P. (1999). Structural basis of DNA recognition by the heterodimeric cell cycle transcription factor E2F-DP. *Genes Dev.* 13, 666–674.
- Zhou, T., Yang, L., Lu, Y., Dror, I., Dantas Machado, A.C., Ghane, T., Di Felice, R., and Rohs, R. (2013). DNASHape: a method for the high-throughput prediction of DNA structural features on a genomic scale. *Nucleic Acids Res.* 41, W56–W62.
- Zhou, T., Shen, N., Yang, L., Abe, N., Horton, J., Mann, R.S., Bussemaker, H.J., Gordân, R., and Rohs, R. (2015). Quantitative modeling of transcription factor binding specificities using DNA shape. *Proc. Natl. Acad. Sci. USA* 112, 4654–4659.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
Bedtools v.2.17	Quinlan and Hall, 2010	http://bedtools.readthedocs.io/
BiasAway	Worsley Hunt et al., 2014	https://github.com/wassermanlab/BiasAway
Biopython v.1.65	Cock et al., 2009	http://biopython.org/
bwtool	Pohl and Beato, 2014	https://github.com/CRG-Barcelona/bwtool
DiMO	Patel and Stormo, 2014	http://stormo.wustl.edu/DiMO
DNASHapeR	Chiu et al., 2016	https://github.com/TsuPeiChiu/DNASHapeR/tree/2nd-order
scikit-learn v.0.16	Pedregosa et al., 2011	http://scikit-learn.org
TFFM	Mathelier and Wasserman, 2013	https://github.com/wassermanlab/TFFM
Other		
DNASHape	Zhou et al., 2013	http://rohslab.cmb.usc.edu/DNASHape/
DNASHapeR	Chiu et al., 2016	https://www.bioconductor.org/packages/release/bioc/html/DNASHapeR.html
GBshape	Chiu et al., 2015	http://rohslab.cmb.usc.edu/GBshape/
Human ChIP-seq peak regions	UCSC genome browser	http://hgdownload.cse.ucsc.edu/goldenPath/hg19/encodeDCC/wgEncodeAwgTfbsUniform/
JASPAR	Mathelier et al., 2014	http://jaspar.genereg.net
Plant ChIP-seq peak regions	Heyndrickx et al., 2014	http://bioinformatics.psb.ugent.be/cig_data/RegNet/

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests will be fulfilled by the corresponding authors Remo Rohs (rohs@usc.edu) and Wyeth W. Wasserman (wyeth@cmmt.ubc.ca).

METHOD DETAILS

ChIP-Seq Datasets and TF Binding Profiles

We retrieved uniformly processed human ENCODE ChIP-seq datasets (Dunham et al., 2012) as narrowPeak-formatted files from the UCSC genome browser at <http://hgdownload.cse.ucsc.edu/goldenPath/hg19/encodeDCC/wgEncodeAwgTfbsUniform/> (as of April 12th 2013). We associated JASPAR TF binding profiles (Mathelier et al., 2014) with ChIPed TFs wherever possible. Using this approach, we obtained 400 ChIP-seq datasets associated with 76 JASPAR profiles (Data S1).

We retrieved *A. thaliana* MADS-domain TF ChIP-seq peak positions studied in Heyndrickx et al. (2014) from the bed-formatted file at http://bioinformatics.psb.ugent.be/cig_data/RegNet/. We specifically considered seven MADS-box ChIP-seq datasets for which a JASPAR TF binding profile was available (Data S1).

ChIP-Seq Peak Sequences

For human ENCODE ChIP-seq peaks, we analyzed 50-bp regions on each side of the peak maximum provided in the narrowPeak-formatted files. Sequences were extracted using the *getfasta* subcommand of bedtools (Quinlan and Hall, 2010) from the hg19 version of the human genome from Ensembl release 63 (Cunningham et al., 2015).

For plant TFs, we considered 50 bp on each side of the middle point of the ChIP-seq peaks because the peak maximum position was not provided. Sequences were extracted using the *getfasta* subcommand of bedtools (Quinlan and Hall, 2010) from the TAIR10 version of the *A. thaliana* genome from Ensembl plant release 26 (Cunningham et al., 2015).

Recurrent ChIP-Seq Peak Regions

We obtained the plots in Data S4 using ChIP-seq peaks found in recurrently ChIPed genomic regions between multiple ChIP-seq experiments for the same TF. Using bedtools, we merged genomic regions where at least two ChIP-seq peaks overlapped among all ChIP-seq experiments associated with the TF. For each merged region, we randomly selected one of the overlapping ChIP-seq peaks. The corresponding set of ChIP-seq peaks was used as the set of foreground sequences in the 10-fold CV.

10-fold CV Datasets

For each ChIP-seq dataset, we constructed 10 training (T_i for $i \in [0, 9]$) and 10 testing (F_i for $i \in [0, 9]$) datasets, where T_i is 9 times the size of F_i . For each training or testing set, we constructed a background dataset (Bt_i or Bf_i , respectively) using the BiasAway tool (Worsley Hunt et al., 2014). Background sequences were obtained by two methods: (i) by randomly selecting the same number of sequences as in T_i or F_i and matching the same %GC composition distribution from a set of genomic background sequences; and (ii) by shuffling sequences in T_i or F_i while matching the same dinucleotide composition.

Genomic background sequences associated with the human ENCODE ChIP-seq datasets were retrieved from mappable regions of the human genome derived from the ENCODE CrgMappability 36-mer track (<ftp://hgdownload.cse.ucsc.edu/goldenPath/hg19/encodeDCC/wgEncodeMapability/wgEncodeCrgMapabilityAlign36mer.bigWig>). We retained mappable regions > 200 bp, which were subsequently split into 100-bp segments.

Genomic background sequences associated with plant ChIP-seq datasets were selected from 1,000,000 random sequences with lengths matched to the considered ChIP-seq regions. Plant sequences were extracted from the TAIR10 genome sequence using the *random* subcommand of bedtools (Quinlan and Hall, 2010).

Discriminative Power Assessment

PSSM- and TFFM-based Model Assessments

We assessed the capacity of models to discriminate ChIP-seq peak regions from background sequences using a 10-fold CV methodology with datasets T_i , F_i , Bt_i , and Bf_i . The PSSM + DNA shape approach was assessed as follows. For all matching training (T_i and Bt_i) and testing (F_i and Bf_i) datasets, each composed of ChIP-seq/foreground and background sequences, we:

1. Optimized the JASPAR binding profile using the DiMO tool (Patel and Stormo, 2014) on the training sequences;
2. Constructed PSSM from the DiMO-optimized profile using the *motifs* Biopython module;
3. Applied PSSM to training sequences (from T_i and Bt_i), and extracted eight DNA shape features at each best site (i.e., with highest PSSM score) per training sequence;
4. Constructed, for each site, a vector combining PSSM score with normalized values of the eight DNA shape parameters at each position;
5. Trained a gradient boosting classifier using the Python *scikit-learn* module (Pedregosa et al., 2011) with vectors associated with the ChIP-seq (T_i) and background (Bt_i) sequences;
6. Applied PSSM to testing sequences (from F_i and Bf_i), and extracted the eight DNA shape features at each best site obtained with the PSSM scores;
7. Constructed corresponding feature vectors, as in 4;
8. Applied the classifier to vectors to obtain the probability of the sequence being a foreground sequence for each testing sequence; and
9. Combined all testing sequence probabilities, and computed AUPRC and AUROC values using the Python *scikit-learn* module (Pedregosa et al., 2011).

To evaluate the TFFM + DNA shape approach on the same sets of training and testing sequences, we:

1. Initialized a TFFM with the DiMO-optimized profile (described above) using the Python TFFM framework (Mathelier and Wasserman, 2013);
2. Trained the TFFM on foreground training sequences (from T_i); and
3. Applied steps 3-9 above using scores from the trained TFFM in place of PSSM.

Predictive powers of the PSSMs and TFFMs were obtained using the same 10-fold CV strategy by solely considering the DiMO-optimized PSSM and trained TFFM scores of the best hit in each sequence.

4-bits-based Classifier Assessment

To evaluate the 4-bits + DNA shape classifiers using the testing and training sets described above, we:

1. Applied the DiMO-optimized PSSM to training sequences (from T_i and Bt_i), and extracted DNA sequence and eight DNA shape features at each best site per training sequence;
2. Constructed, for each site, a vector combining the encoded DNA sequence (A is encoded as 1000, T as 0100, G as 0010, and C as 0001) with the normalized values of the eight DNA shape parameters at each position (Zhou et al., 2015);
3. Trained a gradient boosting classifier using the Python *scikit-learn* module (Pedregosa et al., 2011) with vectors associated with the ChIP-seq (T_i) and background (Bt_i) sequences;
4. Applied the DiMO-trained PSSM to the testing sequences (from F_i and Bf_i), and extracted DNA sequence and eight DNA shape features at each best site obtained with the PSSM scores;
5. Constructed the corresponding feature vectors, as in 2;

6. Applied the classifier to vectors to obtain the probability of the sequence being a foreground sequence for each testing sequence; and
7. Combined all testing sequence probabilities, and computed AUPRC and AUROC values using the Python *scikit-learn* module (Pedregosa et al., 2011).

We obtained predictive powers of the 4-bits classifiers using the same 10-fold CV strategy by solely considering the encoded DNA sequence at the best hit in each ChIP-seq region.

DNA Shape Features

We retrieved values for four DNA shape features (HelT, MGW, ProT, and Roll) and their corresponding second-order shape features (Zhou et al., 2015) for *Homo sapiens* and *A. thaliana* genomes from the GBshape genome browser (Chiu et al., 2015), considering 10-fold CV using the %GC-matched background sequences as genomic sequences. Feature values at each best site in the training and testing sequences were extracted from the corresponding bigWig files using the *extract* subcommand of the *bwtool* software (Pohl and Beato, 2014). For 10-fold CV using the dinucleotide-matched background sequences, DNA shape features were computed with the DNashapeR tool (Chiu et al., 2016) using a customized second-order shape feature branch of the method, provided at <https://github.com/TsuPeiChiu/DNashapeR/tree/2nd-order>. Values of the four DNA shape features and corresponding second-order features were normalized independently by the equation $norm_{value} = (value - min_{value}) / (max_{value} - min_{value})$, where $norm_{value}$ is the normalized value to compute, $value$ is the DNA shape feature value, and min_{value} (max_{value}) corresponds to the minimum (maximum) possible value for the DNA shape feature.

Gradient Boosting Classifiers

We used the GradientBoostingClassifier class implemented in the Python *scikit-learn* module (Pedregosa et al., 2011) to construct, train, and apply gradient boosting classifiers to DNA sequences. Features used in the classifiers to describe a DNA sequence are vectors composed of a 4-bits encoded DNA sequence, PSSM score, or TFFM score and the eight DNA shape features at each nucleotide position of the DNA sequence.

Scanning DNA Sequences with TFFMs and PSSMs

Position frequency matrices (PFMs) corresponding to JASPAR binding profiles were retrieved from the JASPAR database using the *jaspar* BioPython module (Mathelier et al., 2014). The same module was used to convert the DiMO-optimized PFMs to PSSMs using the default background distribution of nucleotides. The module was used to convert the PFMs from JASPAR to PSSMs using the JASPAR pseudocount computation described in the module and default background distribution of nucleotides when assessing feature importance measures. Corresponding PSSMs were used to scan DNA sequences and extract the best hit per sequence.

When applying a trained TFFM to a DNA sequence, the best hit was retrieved from all positions on both strands using the TFFM framework (Mathelier and Wasserman, 2013). First-order hidden Markov model-based TFFMs were used in this study.

TF Family Annotation

We retrieved the family assignment for each TF dataset from the JASPAR database (Mathelier et al., 2014) using the associated binding profile identifiers. This annotation resulted in the inclusion of 76 TFs representing 24 TF families (Data S1).

Feature Importance Measures

We considered gradient boosting classifiers used in 10-fold CV when considering sequence and first-order shape features. To consider the same PSSM per TF for all associated ChIP-seq datasets, we did not apply the DiMO-optimization step and only considered PSSMs derived from the corresponding JASPAR TF binding profiles. We extracted feature importance measures using the Python *scikit-learn* module (Pedregosa et al., 2011). Each measure was averaged over all classifiers used in the 10-fold CV for all datasets associated with a TF. Feature importance measures represent how discriminative a feature is in the underlying decision trees; they do not correspond to feature weights. These measures were visualized using the heat map function of the Python *seaborn* module [<http://dx.doi.org/10.5281/zenodo.19108>].

QUANTIFICATION AND STATISTICAL ANALYSIS

Wilcoxon Signed-Rank Tests

We assessed the significance for the improvement of predictive power when comparing two models using the Wilcoxon signed-rank tests. The *wilcoxon* function of the *scipy.stats* Python module was used in IPython (Perez and Granger, 2007) to compute p values.

Mann-Whitney U Tests

We assessed enrichments for significant improvement of the discriminative power associated with TF families using one-sided Mann-Whitney U tests. The *wilcox.test* R function was used in IPython (Perez and Granger, 2007) through the Python *ipy2* module (<http://rpy.sourceforge.net/>). We corrected Mann-Whitney U test p values for multiple testing by Bonferroni correction. Significant enrichment was defined when $p < 4.17 \times 10^{-4}$, with Bonferroni correction for desired $\alpha < 0.01$.

DATA AND SOFTWARE AVAILABILITY

Software

All output data associated with the 10-fold CV analyses, as well as the IPython notebooks (Perez and Granger, 2007) used to produce associated figures and compute Wilcoxon signed-rank and Mann–Whitney U test p values are provided at https://github.com/amathelier/DNAshapedTFBS_notebooks. Notebooks can be launched by using binder (mybinder.org) at http://mybinder.org/repo/amathelier/DNAshapedTFBS_notebooks. The script *feature_importance_heatmap.py* producing a heat map associated with a single classifier or set of classifiers is provided at <http://github.com/amathelier/DNAshapedTFBS>.

The DNAshaped Python module to train and apply the classifiers for ChIP-seq data can be found at <http://github.com/amathelier/DNAshapedTFBS>.