

Article

ZINCBINDPREDICT — PREDICTION OF ZINC BINDING SITES IN PROTEINS

Sam M. Ireland¹, and Andrew C.R. Martin^{1*} 

¹ Institute of Structural and Molecular Biology, Division of Biosciences, University College London, Darwin Building, Gower Street, London WC1E 6BT, UK; sam.ireland.09@ucl.ac.uk, andrew.martin@ucl.ac.uk

* Correspondence: andrew.martin@ucl.ac.uk

Version January 26, 2021 submitted to *Molecules*

Abstract:

Background: Zinc binding proteins make up a significant proportion of the proteomes of most organisms and, within those proteins, zinc performs rôles in catalysis and structure stabilisation. Identifying the ability to bind zinc in a novel protein can offer insights into its functions and the mechanism by which it carries out those functions. Computational means of doing so are faster than spectroscopic means, allowing for searching at much greater speeds and scales, and thereby guiding complimentary experimental approaches. Typically computational models of zinc binding predict zinc binding for individual residues rather than as a single binding site, and typically do not distinguish between different classes of binding site — missing crucial properties indicative of zinc binding.

Methods: Previously we created ZincBindDB, a continuously updated database of known zinc binding sites, categorised by family (the set of liganding residues). Here we use this dataset to create ZincBindPredict, a set of machine learning methods to predict the most common zinc binding site families for both structure and sequence.

Results: The models all achieve an $MCC \geq 0.88$, recall ≥ 0.93 and precision ≥ 0.91 for the structural models (mean $MCC=0.97$), while the sequence models have $MCC \geq 0.64$, recall ≥ 0.80 and precision ≥ 0.83 (mean $MCC=0.87$), with the models for binding sites containing four liganding residues performing much better than this.

Conclusions: The predictors are available online via a web interface and a GraphQL API at <https://zincbind.bioinf.org.uk/predict/>

Keywords: zinc; metal binding; proteins; prediction; machine learning

0. Introduction

Many proteins require a cofactor to function correctly, and present a region of their surface which has an affinity for that cofactor. Of the metallic cofactors, zinc is one of the most common. Approximately 10% of proteins require zinc to function [1] and so have at least one zinc binding site, making it the second-most prevalent metal in biological systems, after iron. In proteins, it typically performs either a rôle in catalysis (despite, or more likely because of, its lack of variable redox states), or in stabilising a region of the protein [2].

While there are many proteins which are known to bind zinc because the full three-dimensional structure of the protein has been solved in the presence of zinc, leading to the identification of a zinc binding site, it would be useful to be able to determine whether a protein binds zinc without needing to do this. There are experimental means of doing so, but computational approaches offer a more convenient means of performing initial searches at greater scale and speed. These would take either the protein's sequence, or a structure of some kind (either a hypothetical model, an experimental

35 structure generated in the absence of zinc, or an experimental structure solved at low resolution where
36 a zinc cannot be identified, perhaps because of the presence of heavy metals used for isomorphous
37 replacement) and try to predict whether the protein binds zinc, and if so where.

38 There have been numerous studies in this area in the past. Early attempts at predicting zinc
39 binding from sequence were largely done manually, such as by identifying the ‘C...C...H...H’
40 (cys-cys-his-his) motif as being a characteristic indicator of zinc binding [3,4], or by identifying
41 approximate spacing patterns typical of catalytic binding sites — the so-called ‘short and long
42 spacers’ [5]. As the number of available sequences grew and this manual approach became infeasible,
43 sequence alignment with known zinc binding proteins became a useful tool for discovering new
44 zinc binding sites [6,7]. Resources such as PROSITE [8] provide a refinement of manual motif
45 searching by providing motifs for zinc binding in a number of homologous families. At the time
46 of writing, there are 70 motifs for zinc fingers, one for zinc-containing alcohol dehydrogenases, two
47 for copper/zinc superoxide dismutase signature, two for zinc carboxypeptidases and one for the zinc
48 import ATP-binding protein znuC family.

49 By the early 2000s, machine learning became the typical approach for identifying possible metal
50 binding sites — a collection of algorithms which are trained on a dataset of known zinc binding sites
51 in order to identify for themselves what the characteristic properties of zinc binding are, rather than
52 having a human manually identify what those properties might be. Typical algorithms used in the
53 past include Support Vector Machines (SVMs) [9–11] and Random Forests [12,13]. In recent years,
54 deep learning, which relies on multi-layer neural networks to represent the inputs at multiple layers
55 of abstraction, has been used more widely [14,15].

56 Predicting zinc binding from structure has proceeded in a similar fashion, although the nature
57 of structural data means that it has taken longer for there to be enough data to justify the use of
58 machine learning techniques. Early efforts relied on human-observed characteristics of zinc binding
59 sites, such as the ‘hydrophobicity contrast function’, which used the fact that metal binding sites
60 tend to be composed of an inner shell of hydrophilic atoms such as nitrogen and sulphur, which was
61 in turn surrounded by a stabilising shell of hydrophobic atoms [16,17]. As the number of available
62 structures grew, geometric patterns were also observed — both by humans and by machine learning
63 algorithms [17–20]. As with the sequence prediction models, the complexity of the algorithms, and of
64 the zinc binding site features, has grown with the increase in available training data.

65 One recurring feature, particularly in the sequence-based predictive models, is the focus on zinc
66 binding *residues* rather than zinc binding *sites*. In most cases, the entity examined by the predictive
67 model is the individual residue, often with a surrounding linear sequence ‘window’ of residues. The
68 model then assigns a probability as to whether that residue is a zinc binding residue. As outlined
69 above, this approach has had a measure of success, but it is a somewhat artificial concept. There is,
70 after all, no such thing as a zinc-binding residue in isolation. The individual residues of a high-affinity
71 zinc binding site of the kind considered here are only zinc-binding when the other residues are
72 present, and conversely many non-zinc-binding residues could bind zinc if other residues were
73 present in the correct locations. It is particular *combinations* of residues, not individual residues, which
74 are zinc binding — an important fact not usually considered in research of this kind.

75 Another commonality is the treatment of zinc binding sites as a single category, and the
76 presumption of properties that are common to them all regardless of the residues of which they are
77 comprised. This may well be sufficient — particularly as there are essentially only four residues that
78 make up the vast majority of zinc binding sites — but it is possible that properties used for prediction
79 have much tighter distributions within particular sub-categories of zinc binding sites.

80 Previously, we created ZincBindDB [21], a database of zinc binding sites. This resource
81 continuously collates all zinc atoms found in the Protein Data Bank [22], identifies their binding sites
82 (where appropriate), and stores them in a centralised database along with useful properties such as
83 their protein sequence and how different sites cluster together. Sites are classified into ‘families’, not
84 based on homology, but based on the residue composition of the site — the C4 family contains binding

The screenshot displays the ZincBind web interface. On the left, the 'Predict Zinc Binding' panel is active, showing two sub-sections: 'Sequence Prediction' with a text input field for a protein sequence and a 'Predict' button, and 'Structure Prediction' with an 'Upload Structure File (.pdb, .cif or .mmtf)' button and a 'Predict' button. Both sections include a 'Limit families...' dropdown menu. On the right, the 'Sequence Job' panel shows the status 'complete'. It lists the sequence: ETLNNGITDMLTELAFENKIVSQAIRKYNAYRKAASVIAKYPHKIKSGAEAKLPGVGTKIAEKIDEFLATGKLRLEKIRDOOTSSENFLTRVSGIGPSAARKFYDEIGKILEDEGRINEDKLNHQRIGLKYFGDFEKRIPREDLQNDIVLAEVKKVSEYIATVCGSFRGAESSGDMVILLTPSFTSESTKPKLLMOWEQLQKVHFITDLSKGETKFMVGCPLSKMDEK EYPHRRIDILRPLKQDYCYGLVYFGSDIFNKNRMALEKGFINEYTRPLGVTGAV EPLPVDSEKIDFYIQNKYR EPKORSE. It also indicates 'Rejected sites: 2388' and 'Predicted sites: 1'. A note states: 'Note: no predictive model is perfect - predictions here are not guarantees of zinc binding. Some may be false positives, and some false negatives may be missing.' Below the note, the predicted site is shown in bold: **EH2** 0-1. The sequence is displayed with residue numbers 10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150, 160, 170, 180, 190, 200, 210, 220, 230, 240, 250, 260, 270, 280, 290, 300, 310, 320, 330, 340, 350, 360, 370, 380, 390, 400.

Figure 1. The graphical interface for the predictors is shown on the left. The user can enter a protein sequence, or upload a structure file. In both cases the user has the option of limiting the zinc binding families for which the predictor will search which can save a considerable amount of time. Results of the prediction are shown on the right with the residues predicted to form a binding site shown in bold. This interface consumes the ZincBindPredict GraphQL API, which is also publicly available.

85 sites with four cysteines, H3 those with three histidines, and so on. These data are available over the
 86 web via a web 'application programming interface' (API), and using a web interface which provides
 87 three dimensional graphical representations of all the binding sites. As of July 2020, there were 35,811
 88 zinc binding sites in ZincBind, originating from 16,635 PDB structures.

89 We have now used this single, definitive dataset of zinc binding sites to train predictive models
 90 of zinc binding. Here we present models which are trained to detect entire zinc binding sites, rather
 91 than just zinc binding residues, and each predictive model is trained to detect a particular family of
 92 zinc binding sites. There are distinct models for sequence and for structure, and predictions can be
 93 made via the ZincBind website, or via the ZincBindPredict GraphQL API.

94 1. Results and Discussion

95 1.1. Deployment

96 The trained predictive models are available via a simple web interface at
 97 <https://zincbind.bioinf.org.uk/predict/> (see Figure 1). This takes a sequence or an uploaded
 98 PDB file and scans it against each of the models reporting whether any of them suggest a zinc
 99 binding site. Alternatively the ZincBindPredict GraphQL API may be accessed directly. A GraphQL
 100 request can be sent with either a protein sequence or protein structure, and a job ID will be returned.
 101 This can then be polled for results as the protein or sequence is searched using each model in turn,
 102 with the identified binding sites returned as a list with the associated probability.

103 1.2. Training

104 For all twenty datasets (sequence and structure sets each with 10 different combinations
 105 of liganding residues), the ratio of positive samples (actual binding sites) to negative samples
 106 (combinations of residues matching a zinc-binding site family, but which are known not to bind zinc)
 107 was approximately 1:1. The dataset sizes ranged from 804 to 15,332 samples for the sequence datasets,
 108 and from 407 to 3232 samples from the structure datasets.

Table 1. Results for structure models, sorted by Matthews Correlation Coefficient (MCC). The two-residue families' performance was lower than the others as there is essentially just the measurements between two centres to perform the classification, but still scored relatively highly. Four-residue sites in particular were found to have very high performance.

Family	Dataset Size	Recall	Precision	F1	MCC
C2H2	702	1.00	1.00	1.00	1.00
C4	2825	1.00	1.00	1.00	1.00
C3H1	3232	1.00	0.99	1.00	0.99
E1H2	1287	1.00	0.99	1.00	0.99
C2H1	506	1.00	0.98	0.99	0.98
H3	3078	1.00	0.98	0.99	0.98
D1H2	982	1.00	0.98	0.99	0.98
C3	407	1.00	0.98	0.99	0.98
D1H1	522	1.00	0.91	0.95	0.91
E1H1	416	0.93	0.95	0.94	0.88
Mean		0.99	0.98	0.99	0.97

Table 2. Results for sequence models, sorted by Matthews Correlation Coefficient (MCC).

Family	Dataset Size	Recall	Precision	F1	MCC
C4	15332	1.00	0.98	0.99	0.98
H3	4524	0.98	0.99	0.98	0.97
C2H2	3715	0.97	0.99	0.98	0.95
C3H1	9158	0.98	0.96	0.97	0.94
E1H2	2574	0.95	0.97	0.96	0.92
D1H2	2406	0.94	0.95	0.94	0.90
C2H1	1926	0.93	0.95	0.94	0.88
C3	2591	0.95	0.89	0.92	0.84
D1H1	804	0.80	0.93	0.86	0.74
E1H1	812	0.81	0.83	0.82	0.61
Mean		0.93	0.94	0.94	0.87

109 1.3. Models

110 Model effectiveness was measured using recall, precision, F1 score, and Matthews Correlation
 111 Coefficient (MCC) for all twenty models (10 structural and 10 sequence).

112 For the structural models, the lowest MCC score was 0.88 (for the E1H1 model). This, and the
 113 D1H1 model (MCC=0.91), relies on the geometry between just two residues, which makes creating a
 114 distinct separation between the two classes somewhat more difficult — though their performance is
 115 still very close behind that of the three- and four-residue family models. The structure models had an
 116 average MCC of 0.97 (see Table 1).

117 The sequence models also had high scores, though were more variable. The four residue sites
 118 in particular had highly conserved patterns of residue spacing and flanking hydrophobicity despite
 119 being from several homologous families. The average MCC for the sequence models was 0.87, with
 120 the lowest MCC being 0.61 for the E1H1 model and 0.74 for the D1H1 model — again the two
 121 two-residue models were somewhat behind the MCC of 0.84 for the C3 model (see Table 2).

122 While the training is affected by dataset size, this does not appear to be a significant limiting
 123 factor for most of the models. Figure 2 shows the model performance (as MCC) for the sequence
 124 and structure models. The performance of the sequence models falls off as the training set size falls
 125 below ~4000 while the performance of the structural models falls off below around 1000 data points.
 126 The lowest three performing structural models were also the lowest three in dataset size (C3, E1H1,
 127 D1H1), but two of these have only two residues so, as discussed above, the performance might not be




Figure 2. Model Performance (MCC) as a function of training set size. Below $\sim 4,000$ rows, performance declines sharply, though above this threshold there ceases to be a strong correlation between performance and training set size.




Figure 3. Learning curves for all 20 models (10 structural and 10 sequence). Each model was trained on increasing subsets of the overall training set using five-fold cross-validation. Sequence models improved with increasing dataset size whereas, above a low threshold, structure models did not improve with more data.

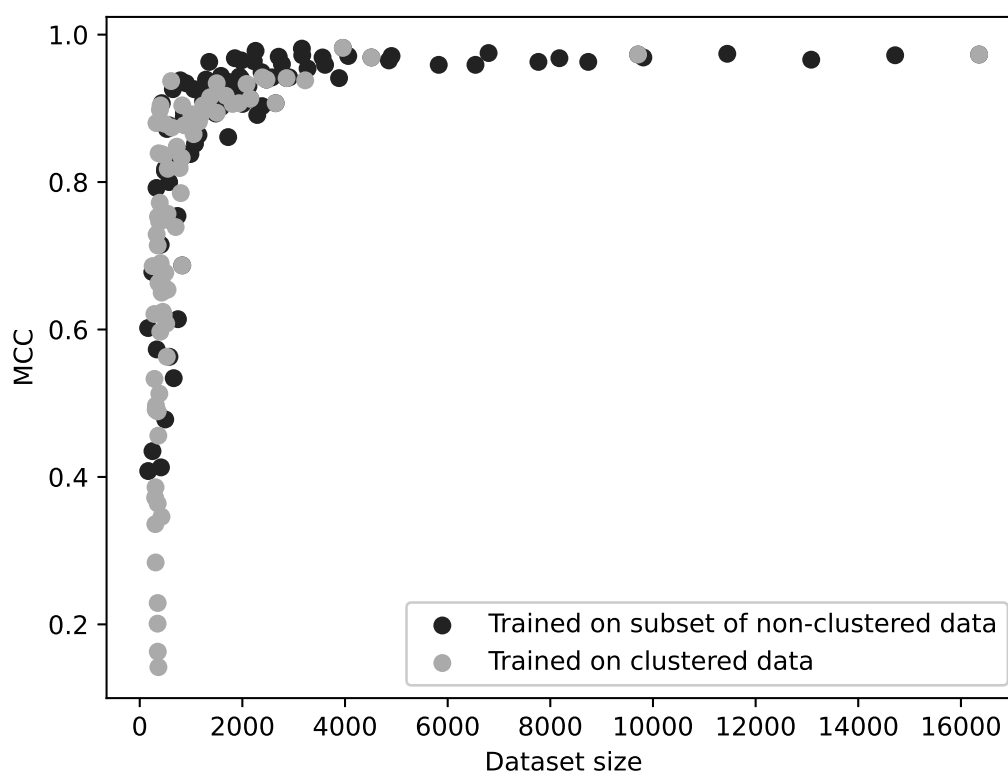


Figure 4. MCC as a function of dataset size for 160 different sequence-based models. For each of the ten zinc-binding site families, 9 classifiers were trained using 20–100% of the original, unclustered data (10×9 models); additional classifiers were trained using sequences clustered at 40–100% sequence identity (10×7 models). The performance (MCC) is plotted against the size of the training dataset. The two modes of dataset reduction are shown by different shades and it can be seen that the curves are not significantly different. This suggests that homology between training and test sets does not influence a model's performance; rather performance is a function of training dataset size.

128 expected to be very good. Learning curves (Figure 3) using fractions of the datasets show a correlation
129 with dataset size for the sequence models, but above around 1000 sequences, the structure models do
130 not improve with larger datasets.

131 The level of abstraction used to describe both sequences and structures (see Methods, Table 5)
132 made it unlikely that any homology between data in the training and testing sets would artificially
133 improve the performance. The features are largely calculated from residues around the binding
134 residues, rather than the sequence in which they occur. Nonetheless, we confirmed that this was
135 true.

136 Different sequence identity thresholds were used for clustering with CD-HIT and, where
137 possible, a dataset of the same size was selected at random from each set of resulting clusters. No
138 significant effect on performance was seen. When clustering at 40% sequence identity, there was
139 slightly lower performance (see Supplementary file `clustering.txt`), but clustering at this level
140 did result in much smaller datasets. As indicated previously, this is a major determinant of the
141 performance of the sequence models.

142 In order to identify whether this lowered performance was because the models performed worse
143 without the possibility of homologous sequences between the training and test sets, or whether it
144 was a result of the smaller training set, for each zinc-binding site family a classifier was trained on
145 20%, 30%, 40%, 50%, 60%, 70%, 80%, 90% and 100% of the original, unclustered data, and additional

Table 3. Predictive ability of BLAST to identify zinc binding sites in protein sequences using homology alone.

Family	Dataset Size	Recall	Precision	F1	MCC
C2H2	3960	0.99	0.95	0.97	0.94
C3H1	9710	0.29	0.87	0.44	0.33
C2H1	2154	0.24	0.88	0.37	0.30
D1H1	818	0.05	0.80	0.09	0.11
C3	2868	0.13	0.61	0.21	0.07
E1H1	828	0.06	0.62	0.11	0.06
D1H2	2470	0.03	0.53	0.06	0.01
H3	5058	0.01	0.19	0.02	-0.10
E1H2	2648	0.02	0.33	0.04	-0.06
Mean		0.18	0.58	0.23	0.17

146 classifiers were trained on data with sequences clustered at 40%, 50%, 60%, 70%, 80%, 90% sequence
147 identity and with no clustering. The performance of the models was then plotted against the resulting
148 dataset sizes as shown in Figure 4. This demonstrates that it is dataset size that determines model
149 performance, regardless of any similarity of the sequences between the training and testing datasets.

150 For reference, the sequence models were compared with using BLAST for predicting
151 zinc-binding sites. For each zinc-binding site family, a BLAST database was created using 80% of
152 the available zinc-binding sequences, and BLAST's ability to identify zinc binding sites from the
153 remaining 20% was compared against an equivalently sized negative set. Results are shown in Table 3.
154 With the exception of C2H2, using BLAST to find zinc binding based on homology performs much
155 worse than the models presented here. Even in the case of C2H2, which seems to have much more
156 similar sequences in its dataset, the ZincBindPredict model still narrowly outperforms BLAST.

157 The performance scores here compare favorably with recent comparable predictive models based
158 on structure and sequence — most notably the 'SVM and Sample-weighted Probabilistic Neural
159 Network' (MCC=0.80) [11], the 'meta-zinc predictor' (MCC=0.79) [23] and ZincExplorer (MCC=0.78)
160 [24].

161 However the models presented here are not intended to be general purpose zinc binding
162 predictors that detect common properties of all zinc binding sites — they are zinc-binding site
163 family-specific predictors based on the principle that common, specific types of zinc binding site
164 have more identifiable, consistent properties than do zinc binding sites in general. As a result, they
165 will not readily detect binding sites of uncommon zinc-binding families. This abstract predictiveness
166 has been deliberately discarded to create highly effective models for specific, common families of
167 zinc binding sites. It is also noteworthy that the binding site itself is a useful unit of prediction using
168 this methodology — even for sequences — rather than individual binding residues. The models are
169 therefore identifying something biologically real (a zinc binding site) rather than something which
170 does not actually exist in isolation (a single zinc binding residue), but which is a useful heuristic in
171 some circumstances.

172 A demonstration of this can be seen by applying the sequence models to bacterial genomes to
173 measure the proportion of typical genomes that the models predict to be zinc binding, as shown for
174 a range of bacterial genomes in Table 4. For most genomes, fewer than 10% of proteins are flagged as
175 zinc binding, with the average for the genomes examined being 8.46%. Given that the zinc-binding
176 families for which predictors have been generated represent 67.0% of binding sites in ZincBindDB
177 (the others being unusual sites), this would imply a 'true' predicted proportion of 12.6% which is a
178 little higher than the widely cited figure of 10%.

Table 4. Percentage of protein sequences encoded in the genome predicted to be zinc binding by ZincBindPredict for an assortment of bacterial genomes. Genomes were acquired from ensembl [25] in the form of translated polypeptide sequences, with a sequence labelled as zinc binding if any of the ten models finds at least one zinc binding site for that sequence/family combination. See Supplementary file *genomes.zip* for the full results.

Species	Percentage of Genome Predicted Zinc Binding
<i>Campylobacter jejuni</i>	6.4%
<i>Clostridioides difficile</i>	5.8%
<i>Enterococcus faecalis</i>	7.5%
<i>Listeria monocytogenes</i>	7.9%
<i>Mycobacterium tuberculosis</i>	11.3%
<i>Salmonella enterica</i>	11.1%
<i>Shigella flexneri</i>	10.1%
<i>Streptococcus pneumoniae</i>	7.6%

Table 5. Details of how features are calculated for residue combinations in structure and sequence models. Hydrophobicity of sequence residues is defined using Wimley and White's scale [26], charge is the count of charged residues (aspartate, glutamate, arginine, histidine and lysine).

Model type	Feature
Sequence	Inter-residue distance (one per gap)
	Average hydrophobicity around residues (window 1)
	Average hydrophobicity around residues (window 3)
	Average hydrophobicity around residues (window 5)
	Average number of charges around residues (window 1)
	Average number of charges around residues (window 3)
	Average number of charges around residues (window 5)
Structure	Mean Inter-C α distance
	Maximum Inter-C α distance
	Minimum Inter-C α distance
	Inter-C α distance standard deviation
	Mean Inter-C β distance
	Maximum Inter-C β distance
	Minimum Inter-C β distance
	Inter-C β distance standard deviation
	Hydrophobic contrast (radius 4 Å)

179 2. Materials and Methods

180 2.1. Dataset creation

181 The datasets used to train the predictive models were derived from ZincBindDB.

182 For the sequence models, for each family of zinc binding sites, all examples were downloaded
183 with the associated sequences, and those with more than one sequence (those sites split across
184 multiple chains) were discarded. The resulting sequences were turned into feature vectors which
185 contained the number of residues between each pair of binding residues, the average hydrophobicity
186 of residues either side of the binding residues, using the features described in Table 5. This created
187 a dataset of positive samples. For the negative samples, for each zinc-binding site family a sequence
188 was chosen at random from the set of all unique sequences in UniProtKB and a combination of
189 residues within that sequence matching the zinc-binding site family (e.g. C2H2), but not a known
190 binding site, was selected — this was done repeatedly until a list of negative samples was built up
191 equal in size to the positive dataset. The two datasets were combined into a single dataset for each
192 zinc-binding site family.

193 For the structural data, for each zinc-binding site family, all relevant zinc binding sites belonging
194 to a PDB structure with resolution better than 2 Ångströms were downloaded, and grouped by the
195 PDB entry to which they belonged. For each PDB entry, the structure was downloaded and parsed
196 using the Python library atomium [27], assembled into the correct biological assembly, and then
197 each binding site was turned into a feature vector using the features described in Table 5. Since
198 the distances used are all the pairwise combinations of the atoms involved, the number of distances
199 depends on the number of liganding residues: H3 sites will have three inter $C\alpha$ distances, C4 sites
200 will have six, and so on. The ‘hydrophobicity contrast function’ is calculated at the centre of the $C\beta$
201 atoms with a radius of 4 Ångströms as described in the original paper by Yamashita *et al.* [16]. This
202 algorithm is a measure of how much outer atoms in a sphere are more hydrophobic than inner atoms,
203 with higher values previously shown to be associated with centres of metal binding [16,17].

204 For example, given a C2H2 site, in a sequence model there would be three inter-residue gaps for
205 which the number of residues per gap would be used together with the mean hydrophobicity and
206 charge of the 4 interacting residues (i.e. a window of 1), the 4 interacting residues plus one sequence
207 neighbour on each side (window of 3) and the 4 interacting residues plus two sequence neighbours
208 on each side (window of 5). For the structural model, there would be 6 inter- $C\alpha$ and 6 inter- $C\beta$
209 distances, from which the mean, maximum, minimum and SD would be calculated as well as the
210 hydrophobicity contrast function.

211 To generate the negative samples, for each positive sample, a random arrangement of
212 residues matching the zinc-binding site family in question was obtained from a randomly chosen,
213 non-zinc-binding PDB structure, and a feature vector created from that non-binding combination. In
214 this case, only residue combinations that could feasibly form a binding site (those where there are no
215 inter- $C\alpha$ distances greater than 30 Ångströms) were used.

216 While the abstraction of sequence and structure suggests that homology is unlikely to influence
217 the results (i.e. homologues between training and testing sets is unlikely to over-rate the performance),
218 this was tested using datasets with similar sequences removed. CD-HIT [28] was used with sequence
219 identity cutoffs ranging from 100% down to 40% (the lowest identity threshold for the standard
220 version of CD-HIT).

221 2.2. Predictive model training

222 The Random Forest algorithm [29] was used to train the predictive model for each of the 20
223 datasets (a dataset of sequence features and a dataset of structural features for each of the ten zinc
224 binding site families), which provided superior results to K-Nearest Neighbours, and vastly superior
225 results to Support Vector Machines even when the dataset was balanced. Random Forests apply the

226 bagging concept (where multiple models are trained on random sub-samples of the data to avoid
227 over-fitting to the training data) to decision trees (classification algorithms which divide the input
228 space into the categories based on sequential binary splits).

229 The hyper-parameters for each model were selected separately using 5-fold cross validation of
230 the training set. The hyper-parameters explored were the impurity measure (gini *vs.* entropy — the
231 algorithm used to split individual trees at each node), the maximum depth that the component trees
232 could have (4, 6, 8 or no maximum), the number of trees in the forest (10, 100 or 1000), and the
233 means of determining the best number of features at each split (either the square root of the number
234 of features, or the log₂ of the number of features). Once optimal hyper-parameters were identified
235 (determined by which combination produced the best F1 score in the cross-validation), the models
236 were trained with those hyper-parameters using the entire training dataset.

237 For the trained model, the metrics recall (how effective at finding true binding sites the models
238 are), precision (how effective at ignoring non-binding sites they are), the F1 score (the harmonic mean
239 of recall and precision) and Matthews Correlation Coefficient (another summary of the true positives,
240 true negatives, false positives and false negatives generally considered the best overall metric [30])
241 were calculated using the separate test datasets (the test-train split being 20:80). The accuracy metric
242 was not used as it is not relevant for unbalanced datasets. Training was performed using the Python
243 scikit-learn library [31].

244 For performance comparison, homology searching was performed using the NCBI BLAST
245 program [32] using an expectation value threshold of 0.1.

246 3. Conclusions

247 Zinc binding sites can be divided into distinct families based on the residues of which they are
248 comprised. These zinc-binding site families follow a power law distribution, with a small number of
249 families being highly represented. By training models for individual zinc-binding site families rather
250 than for zinc binding sites in general, very high recall and precision levels can be achieved. It is worth
251 noting that a zinc-binding site family is a completely different concept from a homologous family as
252 it is the result of convergent evolution potentially spanning many different homologous families. The
253 high performance suggests that, for both sequence and structure, zinc binding properties are more
254 tightly distributed within zinc binding site families than for zinc binding sites generally. The resulting
255 predictor outperforms other general zinc binding predictors.

256 **Acknowledgments:** This work was supported by a Wellcome Trust PhD Studentship to SMI (reference:
257 203756/Z/16/A).

258 **Conflicts of Interest:** The authors declare no conflict of interest. The funders had no rôle in the design of the
259 study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to
260 publish the results.

261 Abbreviations

262 The following abbreviations are used in this manuscript:

263	API	Application Programming Interface
	SVM	Support Vector Machine
264	PDB	Protein Data Bank
	MCC	Matthews Correlation Coefficient

265 References

- 266 1. Andreini, C.; Banci, L.; Bertini, I.; Rosato, A. Counting the zinc-proteins encoded in the human genome.
267 *Journal of Proteome Research* **2006**, *5*, 196–201.
- 268 2. Vallee, B.L.; Auld, D.S. Zinc coordination, function, and structure of zinc enzymes and other proteins.
269 *Biochemistry* **1990**, *29*, 5647–5659.

- 270 3. Miller, J.; McLachlan, A.; Klug, A. Repetitive zinc-binding domains in the protein transcription factor IIIA
271 from *Xenopus oocytes*. *The EMBO journal* **1985**, *4*, 1609–1614.
- 272 4. Barbosa, M.S.; Lowy, D.R.; Schiller, J.T. Papillomavirus polypeptides E6 and E7 are zinc-binding proteins.
273 *Journal of Virology* **1989**, *63*, 1404–1407.
- 274 5. Vallee, B.L.; Auld, D.S. Short and long spacer sequences and other structural features of zinc binding sites
275 in zinc enzymes. *FEBS Letters* **1989**, *257*, 138–140.
- 276 6. Bishop, W.R.; Kirschmeier, P.; George, S.J.; Cramer, S.P.; Hendrickson, W.A.; others. Identification and
277 characterization of zinc binding sites in protein kinase C. *Science* **1991**, *254*, 1776–1779.
- 278 7. Furukawa, Y.; Lim, C.; Tosha, T.; Yoshida, K.; Hagai, T.; Akiyama, S.; Watanabe, S.; Nakagome, K.; Shiro,
279 Y. Identification of a novel zinc-binding protein, C1orf123, as an interactor with a heavy metal-associated
280 domain. *PLoS One* **2018**, *13*, e0204355.
- 281 8. Sigrist, C.J.A.; de Castro, E.; Cerutti, L.; Cuche, B.A.; Hulo, N.; Bridge, A.; Bougueleret, L.; Xenarios,
282 I. New and continuing developments at PROSITE. *Nucleic Acids Research* **2012**, *41*, D344–D347.
283 doi:10.1093/nar/gks1067.
- 284 9. Lin, H.; Han, L.; Zhang, H.; Zheng, C.; Xie, B.; Cao, Z.W.; Chen, Y.Z. Prediction of the functional class
285 of metal-binding proteins from sequence derived physicochemical properties by support vector machine
286 approach. *BMC Bioinformatics* **2006**, *7*, S13.
- 287 10. Srivastava, A.; Kumar, M. Prediction of zinc binding sites in proteins using sequence derived information.
288 *Journal of Biomolecular Structure and Dynamics* **2018**, pp. 1–11.
- 289 11. Li, H.; Pi, D.; Chen, C.; Li, H. A Novel Prediction Method for Zinc-Binding Sites in Proteins
290 by an Ensemble of SVM and Sample-Weighted Probabilistic Neural Network. *IEEE Access* **2019**,
291 *7*, 186147–186157. doi:10.1109/access.2019.2960374.
- 292 12. Zheng, C.; Wang, M.; Takemoto, K.; Akutsu, T.; Zhang, Z.; Song, J. An Integrative Computational
293 Framework Based on a Two-Step Random Forest Algorithm Improves Prediction of Zinc-Binding Sites
294 in Proteins. *PLoS One* **2012**, *7*, 1–15.
- 295 13. Kumar, S. Prediction of Metal Ion Binding Sites in Proteins from Amino Acid Sequences by Using
296 Simplified Amino Acid Alphabets and Random Forest Model. *Genomics & Informatics* **2017**, *15*, 162–169.
- 297 14. Karimi, M.; Wu, D.; Wang, Z.; Shen, Y. DeepAffinity: interpretable deep learning of compound–protein
298 affinity through unified recurrent and convolutional neural networks. *Bioinformatics* **2019**, *35*, 3329–3338.
- 299 15. Haberal, I.; Ogul, H. Prediction of Protein Metal Binding Sites Using Deep Neural Networks. *Molecular*
300 *Informatics* **2019**, *38*, 1800169, [<https://onlinelibrary.wiley.com/doi/pdf/10.1002/minf.201800169>].
- 301 16. Yamashita, M.M.; Wesson, L.; Eisenman, G.; Eisenberg, D. Where metal ions bind in proteins. *Proceedings*
302 *of the National Academy of Sciences, USA* **1990**, *87*, 5648–5652.
- 303 17. Gregory, D.S.; Martin, A.C.R.; Cheatham, J.C.; Rees, A.R. The prediction and characterization of metal
304 binding sites in proteins. *Protein Engineering, Design and Selection* **1993**, *6*, 29–35.
- 305 18. Wallace, A.C.; Borkakoti, N.; Thornton, J.M. TESS: a geometric hashing algorithm for deriving 3D
306 coordinate templates for searching structural databases. Application to enzyme active sites. *Protein*
307 *Science* **1997**, *6*, 2308–2323.
- 308 19. Zhao, W.; Xu, M.; Liang, Z.; Ding, B.; Niu, L.; Liu, H.; Teng, M. Structure-based de novo prediction of
309 zinc-binding sites in proteins of unknown function. *Bioinformatics* **2011**, *27*, 1262–1268.
- 310 20. Liu, Z.; Wang, Y.; Zhou, C.; Xue, Y.; Zhao, W.; Liu, H. Computationally characterizing and comprehensive
311 analysis of zinc-binding sites in proteins. *Biochimica et Biophysica Acta — Proteins and Proteomics* **2014**,
312 *1844*, 171 – 180.
- 313 21. Ireland, S.M.; Martin, A.C.R. ZincBind—the database of zinc binding sites. *Database* **2019**, 2019.
- 314 22. Berman, H.M. The Protein Data Bank. *Nucleic Acids Research* **2000**, *28*, 235–242. doi:10.1093/nar/28.1.235.
- 315 23. Li, H.; Pi, D.; Liang, Y.; Chen, C.; Liu, Y. Integrative computing method for the prediction of
316 zinc-binding sites in proteins. 2017 International Joint Conference on Neural Networks (IJCNN). IEEE,
317 2017. doi:10.1109/ijcnn.2017.7966264.
- 318 24. Chen, Z.; Wang, Y.; Zhai, Y.F.; Song, J.; Zhang, Z. ZincExplorer: an accurate hybrid method to
319 improve the prediction of zinc-binding sites from protein sequences. *Molecular BioSystems* **2013**, *9*, 2213.
320 doi:10.1039/c3mb70100j.
- 321 25. Yates, A.D.; Achuthan, P.; Akanni, W.; Allen, J.; Allen, J.; Alvarez-Jarreta, J.; Amode, M.R.; Armean, I.M.;
322 Azov, A.G.; Bennett, R.; Bhai, J.; Billis, K.; Boddu, S.; Marugán, J.C.; Cummins, C.; Davidson, C.; Dodiya,

- 323 K.; Fatima, R.; Gall, A.; Giron, C.G.; Gil, L.; Grego, T.; Haggerty, L.; Haskell, E.; Hourlier, T.; Izuogu, O.G.;
324 Janacek, S.H.; Juettemann, T.; Kay, M.; Lavidas, I.; Le, T.; Lemos, D.; Martinez, J.G.; Maurel, T.; McDowall,
325 M.; McMahon, A.; Mohanan, S.; Moore, B.; Nuhn, M.; Oheh, D.N.; Parker, A.; Parton, A.; Patricio, M.;
326 Sakthivel, M.P.; Salam, A.I.A.; Schmitt, B.M.; Schuilenburg, H.; Sheppard, D.; Sycheva, M.; Szuba, M.;
327 Taylor, K.; Thormann, A.; Threadgold, G.; Vullo, A.; Walts, B.; Winterbottom, A.; Zadissa, A.; Chakiachvili,
328 M.; Flint, B.; Frankish, A.; Hunt, S.E.; Iisley, G.; Kostadima, M.; Langridge, N.; Loveland, J.E.; Martin, F.J.;
329 Morales, J.; Mudge, J.M.; Muffato, M.; Perry, E.; Ruffier, M.; Trevanion, S.J.; Cunningham, F.; Howe, K.L.;
330 Zerbino, D.R.; Flicek, P. Ensembl 2020. *Nucleic Acids Research* **2019**. doi:10.1093/nar/gkz966.
- 331 26. Wimley, W.C.; White, S.H. Experimentally determined hydrophobicity scale for proteins at membrane
332 interfaces. *Nature Structural Biology* **1996**, *3*, 842–848. doi:10.1038/nsb1096-842.
- 333 27. Ireland, S.M.; Martin, A.C.R. atomium—a Python structure parser. *Bioinformatics* **2020**, *36*, 2750–2754.
334 doi:10.1093/bioinformatics/btaa072.
- 335 28. Fu, L.; Niu, B.; Zhu, Z.; Wu, S.; Li, W. CD-HIT: accelerated for clustering the next-generation sequencing
336 data. *Bioinformatics* **2012**, *28*, 3150–3152.
- 337 29. Breiman, L. Random Forests. *Machine Learning* **2001**, *45*, 5–32. doi:10.1023/a:1010933404324.
- 338 30. Chicco, D.; Jurman, G. The advantages of the Matthews correlation coefficient (MCC) over F1 score and
339 accuracy in binary classification evaluation. *BMC Genomics* **2020**, *21*, 6. doi:10.1186/s12864-019-6413-7.
- 340 31. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.;
341 Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; Duchesnay, E.
342 Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **2011**, *12*, 2825–2830.
- 343 32. Madeira, F.; mi Park, Y.; Lee, J.; Buso, N.; Gur, T.; Madhusoodanan, N.; Basutkar, P.; Tivey, A.R.N.; Potter,
344 S.C.; Finn, R.D.; Lopez, R. The EMBL-EBI search and sequence analysis tools APIs in 2019. *Nucleic Acids
345 Research* **2019**, *47*, W636–W641. doi:10.1093/nar/gkz268.

346 © 2021 by the authors. Submitted to *Molecules* for possible open access publication
347 under the terms and conditions of the Creative Commons Attribution (CC BY) license
348 (<http://creativecommons.org/licenses/by/4.0/>).