



Sequence analysis

EMBER: multi-label prediction of kinase-substrate phosphorylation events through deep learning

Kathryn E. Kirchoff ¹ and Shawn M. Gomez ^{2,3,4,*}

¹Department of Computer Science, The University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA, ²Joint Department of Biomedical Engineering, The University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA, ³North Carolina State University, Raleigh, NC, USA and ⁴Department of Pharmacology, The University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA

*To whom correspondence should be addressed.

Associate Editor: Alfonso Valencia

Received on March 25, 2020; revised on December 9, 2021; editorial decision on February 5, 2022; accepted on February 9, 2022

Abstract

Motivation: Kinase-catalyzed phosphorylation of proteins forms the backbone of signal transduction within the cell, enabling the coordination of numerous processes such as the cell cycle, apoptosis, and differentiation. Although on the order of 10^5 phosphorylation events have been described, we know the specific kinase performing these functions for <5% of cases. The ability to predict which kinases initiate specific individual phosphorylation events has the potential to greatly enhance the design of downstream experimental studies, while simultaneously creating a preliminary map of the broader phosphorylation network that controls cellular signaling.

Results: We describe *Embedding-based multi-label prediction* of phosphorylation events (EMBER), a deep learning method that integrates kinase phylogenetic information and motif-dissimilarity information into a multi-label classification model for the prediction of kinase–motif phosphorylation events. Unlike previous deep learning methods that perform single-label classification, we restate the task of kinase–motif phosphorylation prediction as a multi-label problem, allowing us to train a single unified model rather than a separate model for each of the 134 kinase families. We utilize a Siamese neural network to generate novel vector representations, or an embedding, of peptide motif sequences, and we compare our novel embedding to a previously proposed peptide embedding. Our motif vector representations are used, along with one-hot encoded motif sequences, as input to a classification neural network while also leveraging kinase phylogenetic relationships into our model via a kinase phylogeny-weighted loss function. Results suggest that this approach holds significant promise for improving the known map of phosphorylation relationships that underlie kinome signaling.

Availability and implementation: The data and code underlying this article are available in a GitHub repository at <https://github.com/gomezlab/EMBER>.

Contact: smgomez@unc.edu

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Phosphorylation is the most abundant post-translational modification of protein structure, affecting from one to two-thirds of eukaryotic proteins. In humans, the number of kinases catalyzing this reaction hints at its importance, with kinases being one of the largest gene families with roughly 520 members distributed among 134 families (Lee *et al.*, 2011; Manning *et al.*, 2002; Vlastaridis *et al.*, 2017). During phosphorylation, a kinase facilitates the addition of a phosphate group at a serine, threonine, tyrosine or histidine residue; though other sites exist. Phosphorylation of a substrate at any of these residues occurs within the context of specific consensus phosphorylation sequences, which we refer to herein as ‘motifs’. Additional substrate binding sequences within the kinase or substrate, as well as protein scaffolds that facilitate structural

orientation and downstream catalysis of the reaction, modify the efficacy of motif phosphorylation. Typically, the net effect of kinase phosphorylation is to switch the downstream target into an ‘on’ or ‘off’ state, enabling the transmission of information throughout the cell. Kinase activity touches nearly all aspects of cellular behavior, and the alteration of kinase behavior underlies many diseases while simultaneously establishing the basis for therapeutic interventions (Alunno *et al.*, 2019; Charras *et al.*, 2020; Collins *et al.*, 2018; Deng *et al.*, 2019; Johnson and Lapadat, 2002; Perera *et al.*, 2014; Tegtmeyer *et al.*, 2017; Wilson *et al.*, 2018).

Although the importance of phosphorylation in cellular information processing and its dysregulation as a driver of disease are well-recognized, the map of kinase–motif phosphorylation interactions is mostly unknown. So, while upwards of 100 000 motifs are known

to be phosphorylated, <5% of these have an associated kinase identified as the catalyzing agent (Needham *et al.*, 2019). This knowledge gap provides a considerable impetus for the development of methods aimed at predicting kinase–motif phosphorylation events that, at a minimum, could help focus experimental efforts.

As a result, a number of computational tools have been developed, spanning a myriad of methodological approaches including random forests (Fan *et al.*, 2014), support vector machines (Huang *et al.*, 2015), logistic regression (Li *et al.*, 2018) and Bayesian decision theory (Xue *et al.*, 2006). Advances in deep learning have similarly spawned new approaches, with two methods recently described. The first, MusiteDeep, utilizes a convolutional neural network (CNN) with attention to generate single predictions (Wang *et al.*, 2017). The second deep learning method, DeepPhos, exploits densely connected CNN blocks for its predictions (Luo *et al.*, 2019). Both of these approaches train individual models for each kinase family, requiring a separate model for each of the 134 kinase families. In addition to the practical challenge of training many individual models, a further disadvantage of these two deep learning approaches is the potential lost opportunity gains from transfer learning, as models trained independently do not directly incorporate knowledge of motif phosphorylation by kinases from different kinase families.

2 Approach

Herein, we describe *Embedding-based multi-label prediction of phosphorylation events (EMBER)*, a deep learning approach for predicting *multi-label* kinase–motif phosphorylation relationships. In our approach, we utilize a Siamese neural network, modified for our multi-label prediction task, to generate a high-dimensional embedding of motif vectors. We further utilize one-hot encoded motif sequences. These two representations are leveraged together as a dual input into our classifier, improving learning and prediction. We also find that our Siamese embedding generally outperforms a previously proposed protein embedding, ProtVec, which is trained on significantly more data (Asgari and Mofrad, 2015). We further integrate information regarding evolutionary relationships between kinases into our classification loss function, informing predictions in light of the sparsity associated with the data, and we find that incorporating this information improves prediction accuracy. As EMBER utilizes transfer learning across families, we expect that model accuracy will improve more so than other deep learning approaches as more data describing kinase–substrate relationships are collected. Together, these results suggest that EMBER holds significant promise for improving the known map of phosphorylation relationships that underlie the kinome and broader cellular signaling.

3 Materials and methods

3.1 Kinase–motif interaction data

As documented kinase–motif interactions are sparse in relation to the total number of known phosphorylation events, we attempted to maximize the number of examples of such interactions for training. To do this, we integrated multiple databases describing kinase–motif relationships across multiple vertebrate species. Our data were sourced from PhosphoSitePlus, PhosphoNetworks, and Phospho.ELM, all of which are collections of annotated and experimentally verified kinase–motif interactions (Dinkel *et al.*, 2011; Hornbeck *et al.*, 2012; Hu *et al.*, 2014). From these data sources, non-redundant kinase–motif interactions were extracted and integrated into a single set of kinase–motif interaction pairs. We used the standard single-letter amino acid code for representation of amino acids, with an additional ‘X’ symbol to represent an ambiguous amino acid. Here, motifs are defined as peptide sequences composed of a central phosphorylatable amino acid—either serine (S), threonine (T) or tyrosine (Y)—flanked by seven amino acids on either side. Therefore, each motif is a 15-amino acid peptide or ‘15-mer’. As a phosphorylatable amino acid may not have seven flanking amino acids to either side if it is located near the end of a

substrate sequence, we used ‘–’ to represent the absence of an amino acid in order to maintain a consistent motif length of 15 amino acids across all instances.

The resultant collection of kinase–motif interactions, gleaned from the three aforementioned databases, were all deemed ‘positive’ data. Here, a ‘positive’ data sample is defined as a kinase–motif interaction instance that has been experimentally verified, and that, accordingly, may be traced back to one of the three original databases. Each motif in our dataset is associated with at least one positive label, or rather, is known to be phosphorylated by at least one kinase in our set. Conversely, we define ‘negative’ data to be all remaining kinase–motif pairs for which we found no experimentally verified interactions in any of the source databases.

Deep learning models are known to generally require a large number of examples per class in order to achieve adequate performance. Our original dataset was considerably imbalanced in that all kinase families had a very low positive-to-negative label ratio. For example, the TLK kinase family only has nine positive samples (verified TLK–motif interactions) and more than 10 000 negative samples (lack of evidence for a TLK–motif interaction). To maximize our ability to learn from our data, we utilized only kinases that had a relatively large number of experimentally validated motif interactions, reducing the number of kinase–motif interactions to be used as input for our model. This filtering also served to considerably mitigate the label imbalances in our data. After this initial filtering, we were left with 7531 motifs, from which we set aside 853 motifs for the independent test set, leaving 6678 for the training set. Then, we removed any sequences from the training set that met a 60% similarity threshold with any sequence in the test set, based on Hamming distance scores. This process removed 229 motifs from the training set. Lastly, to be able to make comparisons to previous methods (specifically, MusiteDeep and DeepPhos), it was necessary that we remove any motifs in the test set that could not be directly mapped back to a specific parent protein sequence. After this final truncation, we have a test set of 703 motifs and a training set of 6678 motifs. Kinase labels were then grouped into respective kinase families contingent on data collected from the RegPhos (Lee *et al.*, 2011) database, resulting in eight kinase families. Our resulting dataset is comprised of 7381 phosphorylatable motifs and their reaction-associated kinase families (Table 1). Furthermore, our data are multi-label in that a single motif may be phosphorylated by multiple kinases, including those from other families, resulting in a data point with potentially multiple positive labels.

3.2 Motif embeddings

3.2.1 Protvec embedding

We chose to investigate two methods to achieve our motif embedding. First, we considered ProtVec, a learned embedding of amino acids, originally intended for protein function classification (Asgari and Mofrad, 2015). ProtVec is the result of a Word2Vec algorithm trained on a corpus of 546 790 sequences obtained from Swiss-Prot, which were broken up into three amino acid-long subsequences, or

Table 1. Summary of our kinase–motif phosphorylation dataset

Family	Number of kinases	Number of train motifs	Number of test motifs
Akt	3	382	44
CDK	21	752	112
CK2	2	775	88
MAPK	14	1275	138
PIKK	7	497	56
PKA	5	1235	170
PKC	10	1497	215
Src	11	869	91

Notes: We show the number of kinases per family, and we show the number of motifs phosphorylated by each kinase family in the training set and in the test set.

‘3-grams’. As a result of this approach, ProtVec provides a 100-dimensional distributed representation, analogous to a natural language ‘word embedding’, that establishes coordinates for each possible amino acid 3-gram. This results in a 9048×100 matrix of coordinates, one 100-dimensional coordinate for each 3-gram. In a preliminary investigation, we found that averaging the ProtVec coordinates resulted in a higher-quality embedding compared with the original ProtVec coordinates. Comparisons between the two embeddings are provided in [Supplementary Material](#). We averaged the embedding coordinates, per amino acid, in the following fashion:

We define $T = [\text{AAA}, \text{ALA}, \text{LAA}, \dots, \text{unknown}]$, the vector of 9048 amino acid 3-grams provided by the authors of ProtVec. We also define $A = [\text{A}, \text{L}, \text{S}, \dots, -]$, the alphabet comprising the 22 amino acid symbols. We equate ‘-’ to the ‘unknown’ character defined by ProtVec. Then, we compute the matrix of averaged ProtVec coordinates, $C^{(\text{avg})}$, which will be 22×100 dimensions:

$$C^{(\text{avg})} = \begin{bmatrix} c_{0,0} & c_{0,1} & c_{0,2} & \dots & c_{0,99} \\ c_{1,0} & c_{1,1} & c_{1,2} & \dots & c_{1,99} \\ c_{2,2} & c_{2,1} & c_{2,2} & \dots & c_{2,99} \\ \vdots & & \ddots & & \vdots \\ c_{21,0} & c_{21,1} & c_{21,2} & \dots & c_{21,99} \end{bmatrix} \quad (1)$$

We solve for each element of $C^{(\text{avg})}$ based on the values of $C^{(\text{raw})}$, the original (9048×100) ProtVec matrix:

$$c_{ij}^{(\text{avg})} = \frac{1}{|Q_i|} \sum_{k \in Q_i} c_{kj}^{(\text{raw})} \quad (2)$$

where $c_{ij}^{(\text{avg})}$ belongs to $C^{(\text{avg})}$, $c_{ij}^{(\text{raw})}$ belongs to $C^{(\text{raw})}$, and

$$Q_i = \{q : A_i \in T_q\} \quad (3)$$

Note that the original ProtVec matrix was 9048×100 dimensions, thus each j corresponds to the index of one of the 100 original ProtVec dimensions along the second tensor dimension.

3.2.2 Siamese embedding

We aimed to produce a final model, composed of an embedding technique and a classification method that was specific to our motif dataset. To this end, we implemented a Siamese neural network to provide a novel learned representation of our motifs ([Fig. 1](#)). The Siamese neural network is composed of two identical ‘twin’ neural networks, deemed as such due to their identical hyperparameters as

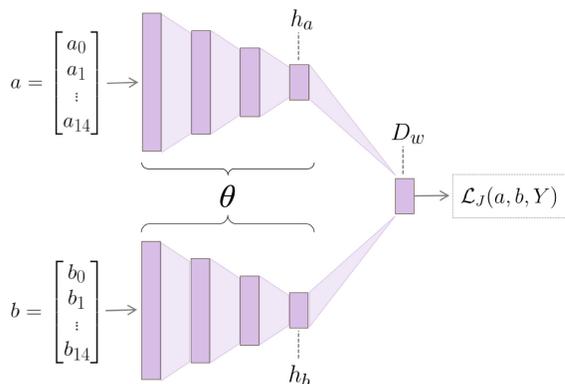


Fig. 1. Siamese neural network architecture, composed of twin CNNs. The twin networks are joined at the final layer. The vectors a and b represent a pair of motifs from the training set, while h_a and h_b represent the respective hidden layers output by either CNN. The difference between the hidden layers is calculated to obtain the distance layer, D_w . D_w is input into the loss along with Y , a variable indicating the dissimilarity, regarding kinase interactions, between a and b . After training is complete, the so-called ‘twin’ architecture is no longer necessary; each motif is input into a single twin and the output of the embedding layer gives the resultant representation of the given motif

well as their identical learned weights and biases ([Bromley et al., 1993](#)). During training, each twin network receives a separate motif sequence that is represented as a one-hot encoding, denoted either as a or b in [Figure 1](#). Motifs are processed through the network until reaching the final fully connected layers, h_a and h_b , which provide the resultant embeddings for the original motif sequences. Next, the layers are joined by calculating the pairwise Euclidean distance, D_w , between h_a and h_b . D_w can be interpreted as the overall dissimilarity between the original motif sequences, a and b . The loss function operates on the final layer, striving to embed relatively more similar data points closer to each other, and relatively more different data points farther away from each other. In this way, the Siamese architecture amplifies the similarities and differences between motifs, and it translates such relationships into a semantically meaningful vector representation for each motif in the embedding space.

We utilized a contrastive loss as described in [Hadsell et al. \(2006\)](#), but we sought to modify the function to account for the multi-label aspect of our task. The canonical Siamese loss between a pair of samples, a and b , is defined as:

$$\mathcal{L}(a, b, Y) = (1 - Y) \frac{1}{2} (D_w)^2 + (Y) \frac{1}{2} [\max(0, m - D_w)]^2, \quad (4)$$

where D_w is the Euclidean distance between the outputs of the embedding layer, m is the margin which is a hyperparameter defined prior to training, and $Y \in \{0, 1\}$. The value of Y is determined by the label of each data point in the pair. If a pair of samples has *identical* labels, they are declared ‘same’ ($Y=0$). Conversely, if a pair of samples has *different* labels, they are declared ‘different’ ($Y=1$). This definition relies on the assumption that each sample may only have one true label. To adapt the original Siamese loss to account for the multi-label aspect of our task, we replaced the discrete variable Y with a continuous variable, namely, the Jaccard distance between kinase-label set pairs. Thus, our modified loss function is defined as:

$$\mathcal{L}_J(a, b, Y) = (1 - J_{a,b}) \frac{1}{2} (D_w)^2 + (J_{a,b}) \frac{1}{2} [\max(0, m - D_w)]^2, \quad (5)$$

where $J_{a,b}$ is shorthand for $J(K_a, K_b)$, which is the Jaccard distance between the kinase-label set K_a and the kinase-label set K_b , associated with motif sample a and motif sample b , respectively. Formally,

$$J(K_a, K_b) = 1 - \frac{|K_a \cap K_b|}{|K_a \cup K_b|} \quad (6)$$

and consequently,

$$0 \leq J(K_a, K_b) \leq 1. \quad (7)$$

In this way, we have defined a continuous metric by which to compare a pair of motifs, rather than the usual ‘0’ or ‘1’ distinction.

The Siamese neural network was trained for 10 000 iterations on the training set, precluding the data points in the independent test set. When composing a mini-batch, we alternated between ‘similar’ and ‘dissimilar’ motif pairs during training. Similar pairs were defined as motifs whose $J(K_a, K_b) > 0.5$, and dissimilar pairs were defined as motifs whose $J(K_a, K_b) \leq 0.5$. After training, we must produce the final embedding space to be used in training of our subsequent classification neural network. To obtain the final embedding, we input each motif into a single arbitrary twin of the original Siamese network (because both twins learn the same weights and biases), producing a high-dimensional (100-dimensional) vector representation of the original motif sequence. The resultant motif embedding effected by the single Siamese twin is further discussed in [Section 4](#). We used k -nearest neighbors (k -NN) classification on each kinase family label to quantitatively compare the predictive capabilities of ProtVec and Siamese embeddings in the coordinate-only space. For our k -NN computation, we used a k of 85.

3.3 Predictive model

3.3.1 Ember architecture

An overview of the architecture of EMBER is shown in [Figure 2](#). EMBER takes as input raw motif sequences and the coordinates of

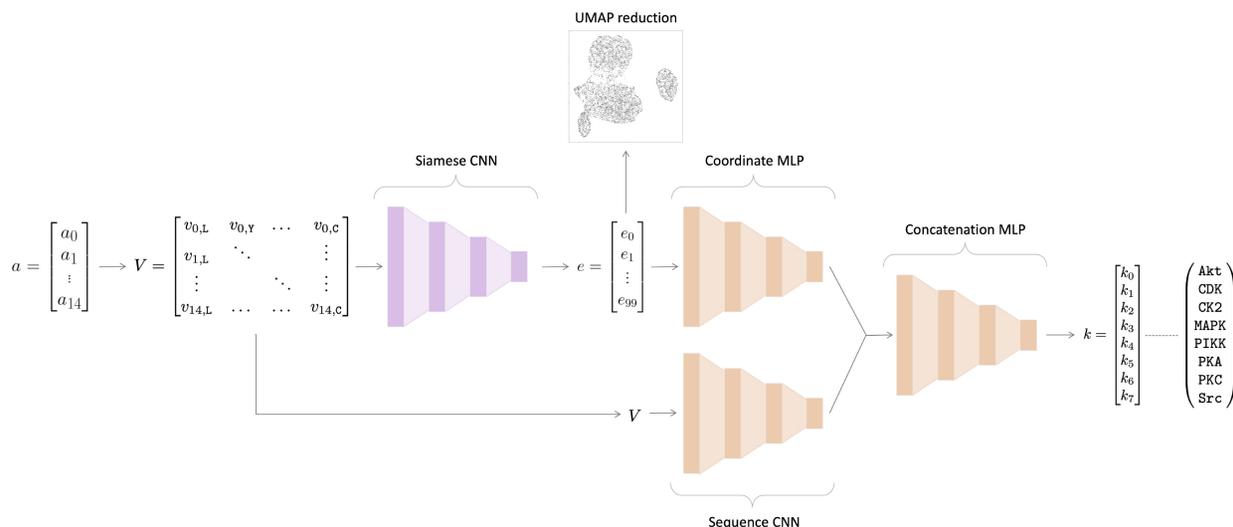


Fig. 2. EMBER model architecture. Here, the previously trained Siamese neural network is colored pink, and the classifier architecture is colored orange. The 15 amino acid-length motif, a , is converted into a one-hot encoded matrix, V . The one-hot encoded matrix is then fed into a single twin from the Siamese neural network. The 100-dimensional embedding, e , is output by the Siamese neural network. Here, we reduce e to a 2D space for illustrative purposes using UMAP. Then, e is fed into a MLP alongside V , which is fed into a CNN. Then, the last layers of the separate neural networks are concatenated, followed by a series of fully connected layers. The final output is a vector, k , of length eight, where each value corresponds to the probability that the motif a was phosphorylated by one of the kinase families indicated in k

each respective motif in the embedding space. We use one-hot encoded motifs as the second type of input into our model. Each motif sequence is represented by a 15×22 matrix. In addition, we utilize the embedding provided by our Siamese neural network, which creates a latent space of dimensions $m \times 90$ where m is the number of motifs.

The inputs into our classifier network, one-hot sequences and embeddings, are fed through a CNN and a multilayer perceptron (MLP), respectively. The outputs of the two networks are then concatenated, and the concatenated layer is fed through a series of fully connected layers (a MLP), followed by a sigmoid activation function. We performed 5-fold cross-validation to assess the accuracy of our model when trained on different training-validation folds. We averaged the performance on the independent test set across the five folds to compute our final performance on the classification task.

3.3.2 Evaluation metrics

In order to quantify the performance of our models, we computed the area under the receiver operating characteristic curve (AUROC) and the area under the precision-recall curve (AUPRC). These metrics were evaluated per kinase family label. In addition to kinase family-specific metrics, we also calculate the micro-average and macro-average for both AUROC and AUPRC. We define $\Lambda = \{\lambda_j : j = 0, \dots, q\}$ as the set of all labels. The micro-average, E_{micro} , aggregates the label-wise contributions of each class:

$$E_{\text{micro}} = E\left(\sum_{\lambda=0}^q tp_{\lambda}, \sum_{\lambda=0}^q tn_{\lambda}, \sum_{\lambda=0}^q fp_{\lambda}, \sum_{\lambda=0}^q fn_{\lambda}\right), \quad (8)$$

where E is an evaluation metric, in our case, either AUROC or AUPRC. Alternatively, the macro-average, E_{macro} , takes into account the score for each respective class and averages those scores together, thus treating all classes equally:

$$E_{\text{macro}} = \frac{1}{q} \sum_{\lambda=0}^q E(tp_{\lambda}, tn_{\lambda}, fp_{\lambda}, fn_{\lambda}), \quad (9)$$

where E is once again an evaluation metric, in our case, either AUROC or AUPRC. Both the E_{macro} and the E_{micro} are calculated based on tp_{λ} , tn_{λ} , fp_{λ} and fn_{λ} , which are, respectively, the number of true positives, the number of true negatives, the number of false positives, and the number of false negatives of label λ .

In order to provide a quantitative comparison between the ProtVec embedding and our Siamese embedding, we calculated silhouette scores for each kinase family label in both embedding spaces (Rousseeuw, 1987). We then took the average across the eight scores for each embedding to get a single silhouette score for each of the two embeddings. Silhouette scores can be used to evaluate the separation quality of clusters in labeled data. For our purposes, a ‘cluster’ refers simply to all motifs phosphorylated by a given kinase family. Reasonably speaking, we would like our clusters to be somewhat well-separated and not overlapping, given that the kinase families we are considering are not all closely related. Calculation of the silhouette score of a cluster depends on two properties: the intracluster distance (average distance between points within the cluster) and the intercluster distance (average distance between the given cluster and neighboring clusters). Silhouette scores range from -1 to 1 , where a value of 1 indicates that the cluster for a given class is relatively far from the remaining clusters and is clearly distinguished, and a value of -1 means the cluster for the given class overlaps with the remaining clusters and is not clearly distinguished.

3.4 Kinase phylogenetic distances

We sought to leverage the phylogenetic relationships between kinases to improve predictions of kinase-motif interactions. Specifically, we considered the dissimilarity of a pair of kinase families in conjunction with the dissimilarity of the two respective groups of motifs that either kinase family phosphorylates (i.e. ‘kinase-family dissimilarity’ vs. ‘motif-group dissimilarity’). Note that the terms ‘distance’ and ‘dissimilarity’ are interchangeable. As the phylogenetic distances given by Manning *et al.* (2002) do not provide distances between typical and atypical kinase families, we established a proxy phylogenetic distance that allows us to define distances between these two types of families. We define this proxy phylogenetic distance through the Levenshtein edit distance, $Lev(k_a, k_b)$, between kinase-domain sequences. Kinase-domain sequences are the specific subsequences of kinases that are directly involved in phosphorylation. These kinase-domain sequences were obtained from an online source provided by Manning *et al.* (2002). Distances between kinase domain sequences were calculated by performing local alignment, utilizing the BLOSUM62 substitution matrix to weight indels and substitutions. To calculate overall kinase-family dissimilarity, we took the average of the Levenshtein edit distances between each kinase domain pair, per family,

$$d(f_a, f_b) = \frac{\sum_{k_a \in f_a} \sum_{k_b \in f_b} Lev(k_a, k_b)}{|f_a| \cdot |f_b|}, \quad (10)$$

where $d(f_a, f_b)$ is the dissimilarity metric (distance) between kinase family a and kinase family b . k_a is the kinase-domain sequence of a kinase belonging to family a , k_b is the kinase-domain sequence of a kinase belonging to family b , and the Levenshtein distance between kinase domain k_a and kinase domain k_b is determined by $Lev(k_a, k_b)$. This formula was applied per kinase family pair and stored in an $a \times b$ kinase-family dissimilarity matrix. From here on, we will refer to this proxy metric for evolutionary dissimilarity between kinase families as the ‘phylogenetic distance’ between kinase families.

3.4.1 Kinase-family dissimilarity versus motif-group dissimilarity

For our (kinase-family dissimilarity)–(motif-group dissimilarity) correlation, we defined motif-group dissimilarity in the same manner as kinase-family dissimilarity, finding the Levenshtein distance between motifs based on local alignment using BLOSUM62. Then, we sought to find the correlation between kinase-family dissimilarity and motif-group dissimilarity. Therefore, calculation of motif-group dissimilarity, per kinase family pair, was defined identically as in Equation (10), but based on the motifs specific to each kinase family, resulting in an $a \times b$ motif-group dissimilarity matrix.

3.4.2 Kinase phylogenetic loss

To leverage evolutionary relationships between kinase families into our predictions, we weighted the original binary cross entropy (BCE) loss by a kinase phylogenetic metric. Specifically, our weighted BCE loss per mini-batch is defined as:

$$PBCE(\hat{y}, y) = -\frac{1}{n} \sum_i P_i^T y_i \log(\hat{y}_i), \quad (11)$$

where n is the size of the mini-batch, y_i is the ground-truth label vector for sample i , \hat{y}_i is the predicted label vector for sample i , and P_i is the phylogenetic weight vector for sample i given by

$$P_i = [w_{0,i}, \dots, w_{|K|,i}]^T, \quad (12)$$

with $w_{k,i}$ being the average phylogenetic weight scalar of label k for sample i :

$$w_{k,i} = \frac{1}{|L_i|} \sum_{j \in L_i} F_{k,j}, \quad (13)$$

and $F_{k,j}$ is the vector of family weights of label k . Finally, L_i is the set of indices corresponding to positive labels for sample i

$$L_i = \{i \in [0, \dots, m-1] : y_i = 1\}, \quad (14)$$

where m is the length of the one-hot true label vector for sample i .

4 Results

4.1 Correlation between kinase phylogenetic dissimilarity and phosphorylated motif dissimilarity

We sought to illuminate the relationship between kinase-family dissimilarity and phosphorylated motif-group dissimilarity described in Section 3. That is, we wanted to determine if ‘similar’ kinases tend to phosphorylate ‘similar’ motifs based on some quantitative metric. To this end, we calculated the correlation between average kinase-family dissimilarities and motif-group dissimilarities based on normalized pairwise alignment scores. From this, we found a Pearson correlation of 0.667, indicating a moderate positive relationship between kinase dissimilarity and that of their respective phosphorylated motifs. Although moderate, this correlation between kinase dissimilarity and motif dissimilarity suggests a potential signal in the phylogenetic relationships that could be leveraged to improve predictions of phosphorylation events.

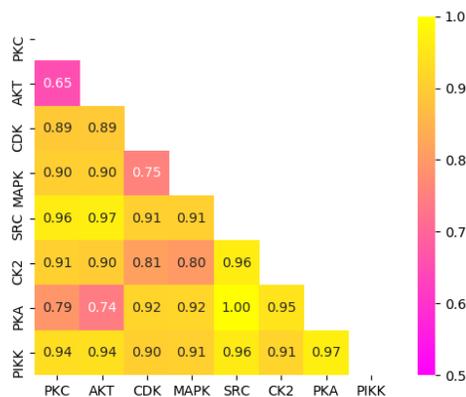


Fig. 3. Heatmap matrix depicting pairwise kinase-domain distances. Levenshtein distances were normalized, with the yellow end of the color bar representing far distances (less similar) and the pink end representing close distances (more similar)

Table 2. AUROC and AUPRC scores on independent test set prediction, given by k -NN performed on the ProtVec and Siamese embeddings

Family	AUROC		AUPRC	
	ProtVec	Siamese	ProtVec	Siamese
Akt	0.908	0.897	0.462	0.513
CDK	0.889	0.892	0.511	0.538
CK2	0.906	0.893	0.665	0.714
MAPK	0.907	0.908	0.739	0.720
PIKK	0.845	0.900	0.579	0.663
PKA	0.865	0.852	0.716	0.659
PKC	0.865	0.885	0.697	0.741
Src	0.998	0.995	0.993	0.991
Macro-average	0.898	0.903	0.670	0.692
Micro-average	0.902	0.924	0.643	0.747

Using our normalized distances as a proxy for phylogenetic distance (see Section 3), the dissimilarity between kinases is displayed as a heatmap in Figure 3. The Akt and PKC families have the greatest similarity (lowest dissimilarity) of all pairwise comparisons, with PKA-Akt and MAPK-CDK following as the next most similar family pairs. Together, these results provide motivation to incorporate both motif dissimilarity and kinase relatedness into the predictive model, as achieved through our custom phylogenetic loss function described in Section 3. The effects of this approach are described in Section 4.

4.2 Motif embedding via Siamese network

We sought to develop a novel learned representation of motifs using a Siamese neural network. Siamese neural networks were first introduced in the early 1990s as a method to solve signature verification, posed as an image-to-image matching problem (Bromley *et al.*, 1993). Siamese neural networks perform metric learning by exploiting the dissimilarity between a pair of data points. Training a Siamese neural network effects a function with the goal of producing a meaningful embedding, capturing semantic similarity in the form of a distance metric. We hypothesized that incorporating high-dimensional vector representations of motifs (i.e. an embedding) into the input of a classification network would provide more predictive power than methods that do not utilize such information. In our Siamese architecture, we opted to use convolutional layers as described in Section 3. We performed k -NN on both the ProtVec and Siamese embeddings of motifs and found that the Siamese embedding produced better predictions, on average, than the ProtVec embedding (see Table 2). More specifically, the Siamese

embedding resulted in a macro-average AUROC of 0.903 compared with ProtVec's 0.898 and a micro-average AUROC of 0.924 compared with ProtVec's 0.902. Likewise, the Siamese embedding had better AUPRC, with a macro-average AUPRC of 0.692 compared with ProtVec's 0.670 and a micro-average AUPRC of 0.747 compared with ProtVec's 0.643. Furthermore, we calculated the silhouette scores of both embeddings and found our Siamese embedding to have a significantly better mean silhouette score of 0.114 compared with ProtVec's 0.005.

We performed dimensionality reduction for visualization of the Siamese embeddings using uniform manifold approximation and projection (UMAP) (McInnes *et al.*, 2018). For our UMAP implementation, we used 200 neighbors, a minimum distance of 0.1 and Euclidean distance for our metric. The resulting 2D UMAP motif embeddings derived from the Siamese neural network are shown in Figure 4. We observe that the motifs phosphorylated by a given

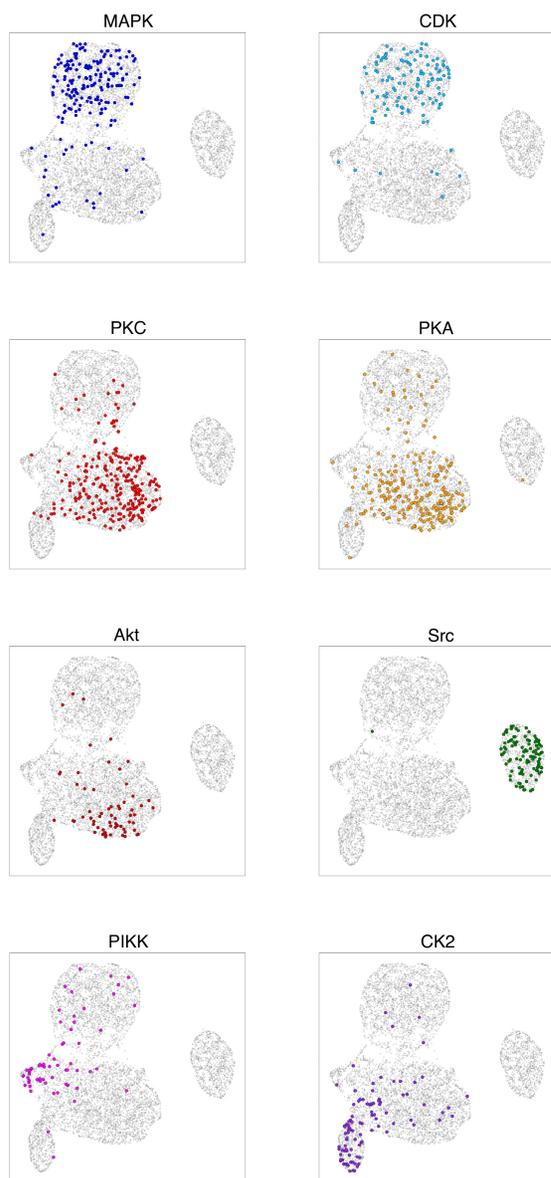


Fig. 4. UMAP projection of the Siamese motif embedding, labeled per kinase family. Each point represents one of the 7302 motifs, and each of the eight panels displays kinase family-specific phosphorylation patterns. Each colored point corresponds to a motif in the test set phosphorylated by a member of the specified kinase family. The axes are the arbitrary x - and y -UMAP dimensions, which are consistent across panels, and are thus not indicated. Furthermore, highlighted points are slightly enlarged to enhance readability

kinase family have a distinctive distribution in the embedding space, with some distributions being highly unique and with some significant overlap between certain families. More specifically, our Siamese embedding shows that motifs phosphorylated by either PKC, PKA, or Akt appear to occupy a similar latent space. In the same vein, motifs phosphorylated by either CDK or MAPK also occupy a similar space. These observations mirror the phylogenetic relationships shown in Figure 3, where the MAPK and CDK families have a relatively short mean evolutionary distance between them, and the PKC–PKA distance, even shorter still.

In addition to these overlapping families, we also observe that Src-phosphorylated motifs form a distinct cluster. This is likely driven by the fact that Src is the only tyrosine kinase family among the eight kinase families we investigated, with its targeted motifs invariably having a tyrosine (Y) at the eighth position in the 15-amino acid sequence, compared with the other 7 families whose motifs have either a serine (S) or a threonine (T) in this position. This effects a significant sequence discrepancy between Src-phosphorylated motifs and remaining motifs. The fact that Src-phosphorylated motifs cluster so precisely serves as a sanity check that our Siamese embedding is capturing sequence (dis)similarity information despite being trained through comparison of kinase–motif phosphorylation events in lieu of explicit motif sequence comparisons. We note that the embedding produced by our Siamese neural network is quite qualitatively similar to the ProtVec embedding in terms of these kinase-label clusters indicated in the UMAP projections. The UMAP projections of the ProtVec embeddings are included in Supplementary Material.

4.3 Prediction of phosphorylation events

Following training of EMBER on both motif sequences and motif vector representations as input, we conducted an ablation test in which we removed the motif vector representation (or coordinate) input along with its respective MLP; this was achieved by applying a dropout rate of 1.00 on the final layer of the coordinate-associated MLP. This ablation test allowed us to observe how our novel motif sequence-coordinate model compares to a canonical deep learning model whose input consists solely of one-hot encoded motif sequences (such as in the methods utilized by Luo *et al.*, 2019; Wang *et al.*, 2017). We also compared EMBER trained with the standard BCE loss to EMBER trained with our kinase phylogenetic loss. All predictive models, as described in Table 3, were trained on identical training-validation splits and evaluated on the same independent test set.

Comparisons between the predictive capability of the models described here are quantified by AUROC and AUPRC, and these metrics are presented for each of the three models in Table 3. As indicated by Table 3, EMBER, utilizing both sequence and coordinate information, outperforms the canonical sequence model in both AUROC and AUPRC. In addition, integration of phylogenetic information into the loss provides a generally small but consistent additional boost in performance, showing the best overall results out of the three models for AUROC and AUPRC. Individual performance metric curves for each kinase label, produced by EMBER trained via the phylogenetic loss, are shown in a subplot of Figure 5.

A confusion matrix providing greater detail and illustrating the relative effectiveness of our model for prediction of different kinase families is shown in Figure 6. In order to compute the confusion matrix, we set a prediction threshold of 0.5, declaring any prediction above 0.5 as ‘positive’ and any prediction equal to or <0.5 as ‘negative’. As indicated by the confusion matrix, the model often confounds motifs that are phosphorylated by closely related kinase families, for example, MAPK and CDK. This is presumably due to the close phylogenetic relationship between MAPK and CDK, as indicated by their relatively low phylogenetic distance of 0.75 (Fig. 3). Furthermore, our Siamese neural network embeds motifs of these respective families into the same relative space, as shown in Figure 4, further illustrating the confounding nature of these motifs. A similar trend is found for motifs phosphorylated by PKC, PKA and Akt. This trio is also shown to be closely related as indicated by the correlations in Figure 3 and the embeddings in Figure 4.

Table 3. AUROC and AUPRC results achieved on the independent test set across deep learning classification models

Family	AUROC			AUPRC		
	Seq-CNN	EMBER (<i>BCE</i>)	EMBER (<i>PBCE</i>)	Seq-CNN	EMBER (<i>BCE</i>)	EMBER (<i>PBCE</i>)
Akt	0.828	0.858	0.881	0.327	0.323	0.423
CDK	0.906	0.917	0.925	0.648	0.696	0.698
CK2	0.900	0.914	0.923	0.722	0.780	0.788
MAPK	0.887	0.899	0.896	0.665	0.694	0.697
PIKK	0.871	0.886	0.913	0.586	0.611	0.670
PKA	0.846	0.868	0.877	0.636	0.680	0.705
PKC	0.873	0.895	0.902	0.735	0.774	0.798
Src	0.997	0.994	0.996	0.992	0.992	0.994
Macro-average	0.889	0.904	0.914	0.664	0.694	0.722
Micro-average	0.913	0.926	0.932	0.731	0.763	0.784

Notes: The AUROC and AUPRC are presented per kinase family for each model. From left to right, we include results for the ablated sequence-only CNN, EMBER trained using a canonical BCE loss, and EMBER trained using the kinase phylogeny-weighted loss as described in Section 3.

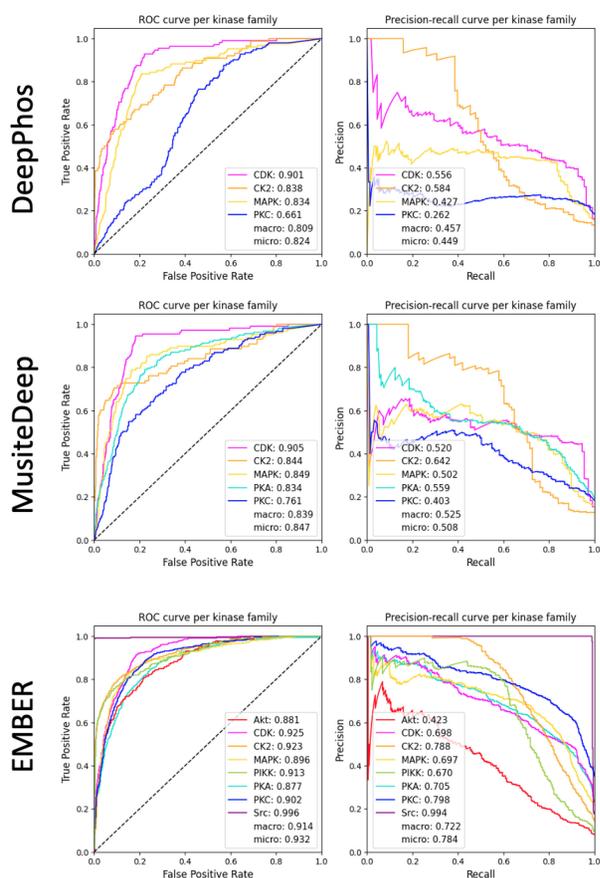


Fig. 5. AUROC and AUPRC results achieved on the independent test set by DeepPhos, MusiteDeep and EMBER. The AUROC and AUPRC of each kinase family label are shown in the respective legends. Note that color-coding is consistent across subplots.

4.3.1 Comparison to existing methods

We sought to compare EMBER's performance to the two existing deep learning methods, MusiteDeep and DeepPhos, which adopt single-label models. However, this is not a straight-forward comparison because EMBER was trained on sequences 15 amino acids in length while MusiteDeep and DeepPhos require sequences of 33 and 51 amino acids in length, respectively. Therefore, in order to make our test set suitable for MusiteDeep and DeepPhos, it is necessary that we unambiguously map each 15-mer subsequence to a complete substrate sequence, from which we can obtain both a

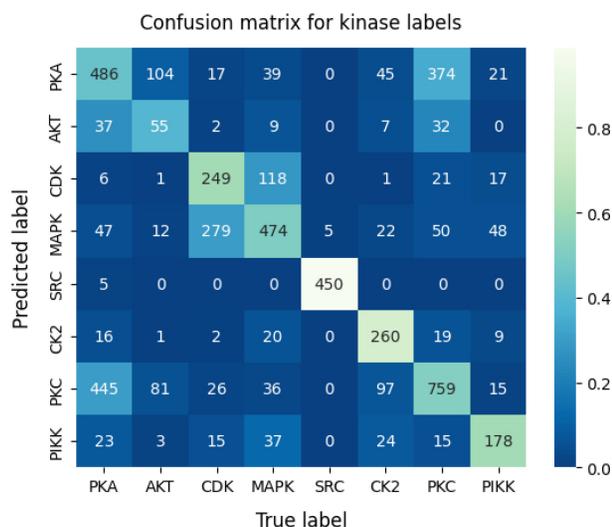


Fig. 6. Confusion matrix for EMBER predictions on the test set. The numbers inside each box represent the raw number of predictions per box. The color scale is based on the ratio of predictions (in the corresponding box) to total predictions, per label. A lighter color corresponds to a larger ratio of predictions to total predictions

33-mer and a 51-mer. This was accomplished by simply referencing the source databases, which either directly provided a complete substrate sequence or provided a 15-mer subsequence and an accession number (via which the associated substrate sequence can be easily retrieved). After obtaining the set of complete substrate sequences associated with our test set, we were able to elongate each motif to 33 and 51 amino acids in length for evaluation through MusiteDeep and DeepPhos, respectively.

We note that of the eight kinase families for which our model produces predictions, DeepPhos has functioning models for only four of the families (CDK, CK2, MAPK and PKC), and MusiteDeep has models for only five of the families (CDK, CK2, MAPK, PKA and PKC). We show AUROC and AUPRC results per kinase label from each of the three methods in Figure 5. EMBER outperforms MusiteDeep and DeepPhos on all four averaged metrics, indicating that our multi-label approach may be better equipped to solve the problem of kinase-motif prediction compared with the single-label approaches.

5 Discussion

Illuminating the map of kinase-substrate interactions has the potential to enhance our understanding of basic cellular signaling as well

as drive health applications, e.g. by facilitating the development of novel kinase inhibitor-based therapies that disrupt kinase signaling pathways. In this work, we have presented a deep learning-based approach that aims to predict which substrates are likely to be phosphorylated by a specific kinase family. In particular, our multi-label approach establishes a unified model that utilizes all available kinase-motif data to learn broader structures within the data and improve predictions across all kinase families in tandem. This approach avoids challenges in hyperparameter tuning inherent in the development of an individual model for each kinase. We believe that this approach will enable continuing improvement in predictions, as newly generated data describing any kinase-motif phosphorylation event can assist in improving predictions for all kinases. That is, a kinase-motif interaction discovered for PKA will improve the predictions not just for PKA, but also for Akt, PKC, MAPK etc. through the transfer learning capabilities inherent in our multi-label model.

We showed that incorporation of a learned vector representation of motifs, namely the motifs' coordinates in the Siamese embedding space, serves to improve performance over a model that utilizes only one-hot encoded motif sequences as input. Not only did the Siamese embedding improve prediction of phosphorylation events through a neural network architecture, but it also outperformed ProtVec, a previously developed embedding, in a coordinate-based k -NN task. This improvement over ProtVec was in spite of the fact that our Siamese neural network utilized <7000 training sequences of 15 amino acids in length compared with ProtVec's 500 000 sequences of ~300 amino acids in average length. The Siamese embedding was further generated through direct comparison of kinase-motif phosphorylation events rather than simply the sequence-derived data used by ProtVec. Furthermore, ProtVec is a generalized protein embedding while the Siamese embedding described here has the potential to be customized. For example, the use of the Jaccard distance in the Siamese loss allows the neural network to be trained on any number of multi-label datasets such as acetylation, methylation or carbonylation reactions. We also found that there is a small though meaningful relationship between the evolutionary distance between kinases and the motifs they phosphorylate, supporting the concept that closely related kinases will tend to phosphorylate similar motifs. When encoded in the form of our phylogenetic loss function, this relationship was able to slightly improve prediction accuracies. Together, these results suggest that EMBER holds significant promise towards the task of illuminating the currently unknown relationships between kinases and the substrates they act on.

Acknowledgements

We would like to acknowledge members of the GomezLab for helpful comments and feedback.

Funding

This work was supported by grants through the National Institutes of Health [grant numbers CA177993, CA233811, CA238475 and DK116204].

Conflict of Interest: none declared.

References

Alunno, A. *et al.* (2019) Pathogenic and therapeutic relevance of JAK/STAT signaling in systemic lupus erythematosus: integration of distinct

- inflammatory pathways and the prospect of their inhibition with an oral agent. *Cells*, **8**, 898.
- Asgari, E. and Mofrad, M.R.K. (2015) Continuous distributed representation of biological sequences for deep proteomics and genomics. *PLoS One*, **10**, e0141287.
- Bromley, J. *et al.* (1993) Signature verification using a "Siamese" time delay neural network. In: Proceedings of the 6th International Conference on Neural Information Processing Systems, NIPS'93, pp. 737–744, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Charras, A. *et al.* (2020) JAK inhibitors suppress innate epigenetic reprogramming: a promise for patients with Sjögren's syndrome. *Clin. Rev. Allergy Immunol.*, **58**, 182–193.
- Collins, K.A.L. *et al.* (2018) Proteomic analysis defines kinase taxonomies specific for subtypes of breast cancer. *Oncotarget*, **9**, 15480–15497.
- Deng, Y.N. *et al.* (2019) Essential kinases and transcriptional regulators and their roles in autoimmunity. *Biomolecules*, **9**, 145.
- Dinkel, H. *et al.* (2011) Phospho.ELM: a database of phosphorylation sites—update. *Nucleic Acids Res.*, **39** (Database issue), D261–D267.
- Fan, W. *et al.* (2014) Prediction of protein kinase-specific phosphorylation sites in hierarchical structure using functional information and random forest. *Amino Acids*, **46**, 1069–1078.
- Hadsell, R. *et al.* (2006) Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), Vol. 2, pp. 1735–1742, New York, NY, USA. IEEE.
- Hornbeck, P.V. *et al.* (2012) PhosphoSitePlus: a comprehensive resource for investigating the structure and function of experimentally determined post-translational modifications in man and mouse. *Nucleic Acids Res.*, **40** (Database issue), D261–D270.
- Hu, J. *et al.* (2014) PhosphoNetworks: a database for human phosphorylation networks. *Bioinformatics*, **30**, 141–142.
- Huang, S.-Y. *et al.* (2015) Using support vector machines to identify protein phosphorylation sites in viruses. *J. Mol. Graph. Model*, **56**, 84–90.
- Johnson, G.L. and Lapadat, R. (2002) Mitogen-activated protein kinase pathways mediated by ERK, JNK, and p38 protein kinases. *Science*, **298**, 1911–1912.
- Lee, T.-Y. *et al.* (2011) RegPhos: a system to explore the protein kinase-substrate phosphorylation network in humans. *Nucleic Acids Res.*, **39** (Database issue), D777–D787.
- Li, F. *et al.* (2018) Quokka: a comprehensive tool for rapid and accurate prediction of kinase family-specific phosphorylation sites in the human proteome. *Bioinformatics*, **34**, 4223–4231.
- Luo, F. *et al.* (2019) DeepPhos: prediction of protein phosphorylation sites with deep learning. *Bioinformatics*, **35**, 2766–2773.
- Manning, G. *et al.* (2002) The protein kinase complement of the human genome. *Science*, **298**, 1912–1934.
- McInnes, L. *et al.* (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. *J. Open Source Softw.*, **3**(29), 861.
- Needham, E.J. *et al.* (2019) Illuminating the dark phosphoproteome. *Sci. Signal.*, **12**, eaau8645.
- Perera, G.K. *et al.* (2014) Integrative biology approach identifies cytokine targeting strategies for psoriasis. *Sci. Transl. Med.*, **6**, 223ra22.
- Rousseeuw, P.J. (1987) Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.*, **20**, 53–65.
- Tegtmeyer, N. *et al.* (2017) Subversion of host kinases: a key network in cellular signaling hijacked by helicobacter pylori CagA. *Mol. Microbiol.*, **105**, 358–372.
- Vlastaridis, P. *et al.* (2017) Estimating the total number of phosphoproteins and phosphorylation sites in eukaryotic proteomes. *Gigascience*, **6**, 1–11.
- Wang, D. *et al.* (2017) MusiteDeep: a deep-learning framework for general and kinase-specific phosphorylation site prediction. *Bioinformatics*, **33**, 3909–3916.
- Wilson, L.J. *et al.* (2018) New perspectives, opportunities, and challenges in exploring the human protein kinome. *Cancer Res.*, **78**, 15–29.
- Xue, Y. *et al.* (2006) PPSP: prediction of PK-specific phosphorylation site with Bayesian decision theory. *BMC Bioinformatics*, **7**, 163.