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

Kathryn E. Kirchoff[1] **and Shawn M. Gomez**[2,3,✉]

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

Abstract

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. While on the order of 10^5 phosphorylation events have been described, we know the specific kinase performing these functions for less than 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. To this end, we describe EMBER, a deep learning method that integrates kinase-phylogeny 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 network to generate novel vector representations, or an embedding, of 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 network while also leveraging kinase phylogenetic relationships into our model via a kinase phylogeny-based loss function. Results suggest that this approach holds significant promise for improving our map of phosphorylation relations that underlie kinome signaling.

Availability: https://github.com/gomezlab/EMBER
Correspondence: *smgomez@unc.edu*

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 (1–3). During phosphorylation, a kinase facilitates the addition of a phosphate group at serine, threonine, tyrosine, or histidine residues, 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 here 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 (4–11).

Although the importance of phosphorylation in cell information processing and its dysregulation as a driver of disease is well recognized, the map of kinase-motif phosphorylation interactions is mostly unknown. So, while upwards of 100,000 motifs are known to be phosphorylated, less than 5% of these have an associated kinase identified as the catalyzing agent (12). 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 (13), support vector machines (14), logistic regression (15), and Bayesian decision theory (16). Advances in deep learning have similarly spawned new approaches, with two methods recently described. The first, MusiteDeep, utilizes a convolutional neural network (CNN) and a subsequent attention layer to generate single predictions (17). The second deep learning method, DeepPhos, exploits densely connected CNN (DC-CNN) blocks for its predictions (18). Both of these approaches train individual models for each kinase family, requiring separate models for each of the 134 families. In addition to this practical challenge, a further disadvantage of these approaches is the potential lost opportunity gains from transfer learning, as models do not directly incorporate motif phosphorylation by kinases from different kinase families.

Here, we describe, EMBER (Embedding-based multi-label prediction of phosphorylation events), 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. Along with 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 this approach performs as well or better than a previously proposed protein embedding, ProtVec, trained on significantly more data (19). We further integrate information regarding evolutionary relationships between kinases into our classification network loss function, informing predictions in light of the sparsity associated with these data, and we find that 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 as data describing kinase-substrate relationships increases. Together, these results suggest that this approach holds significant promise for improving our map of phosphorylation relationships that underlie the kinome and broader cellular signaling.

Methods

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 datasets describing motif-kinase relationships across multiple vertebrate species. Our data was sourced from PhosphoSitePlus, PhosphoNetworks, and Phospho.ELM, all of which are collections of annotated and experimentally verified motif-kinase relationships (20), (21), (22). From these data sources, non-redundant motif-kinase relationships were extracted and integrated into a single set of interactions. 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. We defined our motifs as peptides composed of a central phosphorylatable amino acid — either serine (S), threonine (T), or tyrosine (Y) — flanked by 7 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 7 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.

Deep learning models are known to generally require large amounts of examples per class in order to achieve adequate performance. Our original dataset was considerably imbalanced in that all positive labels (verified kinase-motif interactions) had a very low positive- to negative-label ratio. For example, the TLK kinase family only has 9 positive labels (verified TLK-motif interactions) and more than 10,000 negative labels (lack of evidence for a TLK-motif interaction). To maximize our ability to learn from our data, we only utilized kinases that had a relatively large number of experimentally validated motif interactions as input into the model, reducing the number of kinase-motif relationships to be used for our model. However, this data filtering also served to considerably mitigate the imbalances in our data. Kinases were then grouped into respective kinase families contingent on data collected from the RegPhos (1) database, resulting in eight kinase families. The number of kinases per family is presented in Table 1. Our resulting data set included 7535 phosphorylatable motifs and their reaction-associated kinase families (Table 1). Furthermore, our data is 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. We set aside 857 motifs for the independent test set.

Table 1. Summary of the number of identified motifs phosphorylated by each kinase family.

Kinase family	Number of motifs
Akt	464
CDK	903
CK2	900
MAPK	1514
PIKK	574
PKA	1533
PKC	1801
Src	993

Motif embeddings.

ProtVec embedding. We chose to investigate two methods to achieve our motif embedding. First, we considered ProtVec, a learned embedding of amino acids, originally meant for protein function classification (19). ProtVec is the result of a Word2Vec implementation trained on a corpus of 546,790 sequences obtained from Swiss-Prot, which were broken up into 3 amino acid-long subsequences. As a result, ProtVec provides a 100-dimensional distributed representation, or "word embedding", that establishes coordinates for each possible amino acid 3-gram, resulting in a 9048 x 100 matrix. We averaged the embedding coordinates, per amino acid, reducing the first tensor dimension from 9048 to 22.

Siamese embedding. We implemented a Siamese network to provide a novel learned representation of our motifs (Figure 1). The Siamese network is composed of two identical "twin" networks, deemed as such due to their identical hyperparameters as well as their identical learned weights and biases (23). During training, each twin network receives a separate motif sequence, represented as a one-hot encoding, and 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 between h_a and h_b, resulting in d_a and d_b. The pairwise mean between d_a and d_b can be interpreted as the overall dissimilarity between m_a and m_b, represented by D. 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. In this way, the network 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.

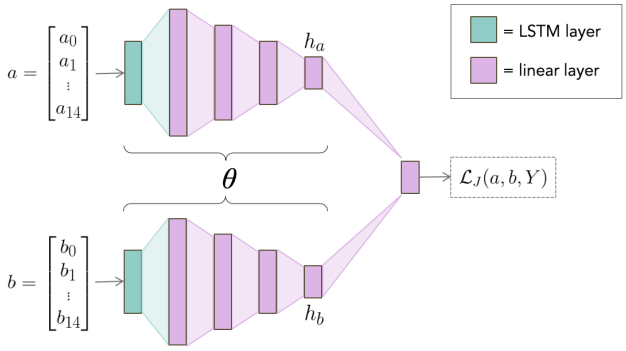

Fig. 1. Siamese network architecture, composed of twin neural networks, each containing identical long short-term memory (LSTM) layers followed by fully-connected layers. The twin networks are joined at the final layer. a and b represent a pair of motifs from the training set, while h_a and h_b represent the respective hidden layers output by either LSTM. The difference between the hidden layers is calculated to obtain the distance layer, which is then averaged to get a final scalar value representing the dissimilarity between the motifs. After training is complete, each motif is input into a single twin and the output of the embedding layer gives the resultant representation of the given motif.

We utilized a contrastive loss as described in (24), but we modified 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

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

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

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

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,

$$J(K_a, K_b) = 1 - \frac{|K_a \cap K_b|}{|K_a \cup K_b|} \quad \textbf{(3)}$$

and consequently,

$$0 \leq J(K_a, K_b) \leq 1. \quad \textbf{(4)}$$

In this way we have defined a continuous metric by which to compare a pair of motifs, rather than the usual 0-1 distinction. The Siamese network was trained for 250,000 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 network. To obtain the final embedding, we input each motif into a single arbitrary twin of the original network (because both twins learn the same weights and biases), producing a high-dimensional (1500-dimensional) vector representation of the original motif sequence. The resultant motif embedding effected by the single Siamese twin is further discussed in the Results section. We used k-nearest neighbors classification on each family to quantitatively compare the predictive capabilities of ProtVec and Siamese embeddings in the coordinate-only space.

Predictive model framework. An overview of the architecture of EMBER is shown in Figure 2. EMBER takes as input raw motif sequences and the coordinates of 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 x 22 vector. In addition, we utilize the embedding provided by our Siamese network, which creates a latent space of dimensions m x 1500 where m is the number of motifs. The space is reshaped to m x 15 x 100 such that each motif coordinate is 15 x 100.
We utilized convolutional neural networks (CNNs) in lieu of recurrent neural networks (RNNs) which are typically used in prediction from sequential data due because our preliminary work showed better performance of the former. Each type of input, both one-hot sequences and embeddings, is introduced to the network through their respective CNNs, followed by downstream fully-connected (FC) layers. The dual CNN-FC networks have identical hyperparameters in an effort to equally weight both sequence and coordinate input. The CNN-FC networks are then concatenated in the final layer, which is then fed through a sigmoid activation function and outputs a 8 x 1 vector per motif.

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. (2), 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 families. We define this proxy phylogenetic distance through the Levenshtein edit distance, $Lev_{(k_a,k_b)}$, between

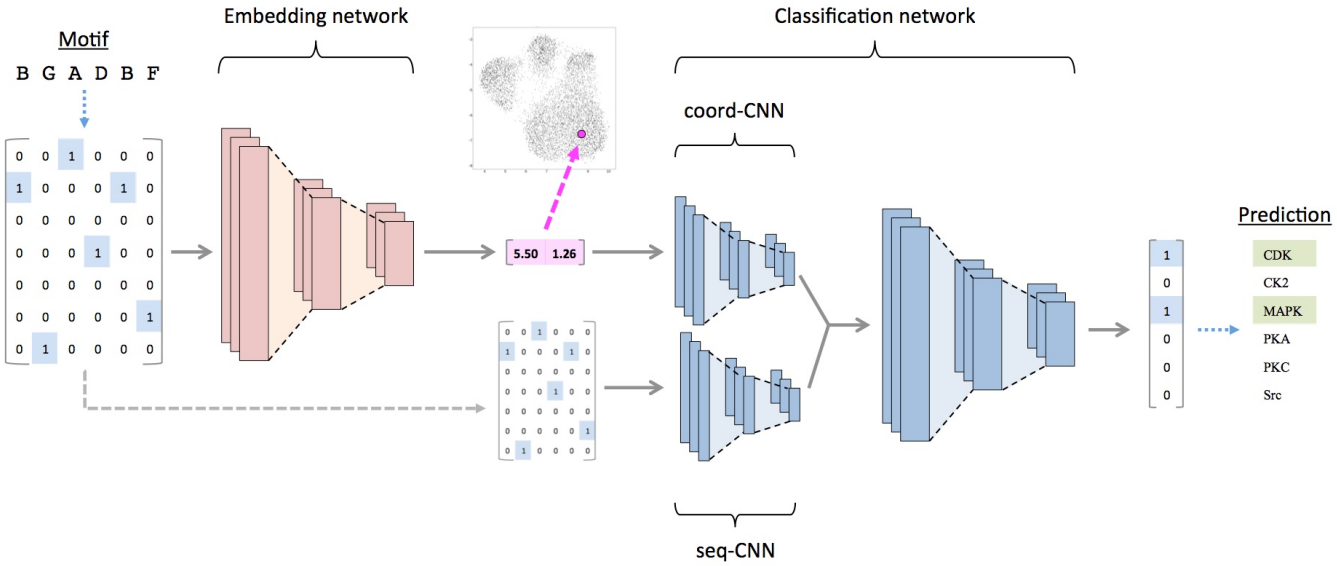

Fig. 2. EMBER model architecture. For simplification, we show a 6-amino acid motif in conjunction with a 7-word vocabulary. Here, the motif is converted into a one-of-seven encoding in order to create a proper input vector of dimensions 7 x 6; the implemented model uses a 22-word vocabulary and motifs of length 15. The one-hot encoded motif is fed into the embedding network, which is a single twin of the pre-trained Siamese network. The figure shows a 2-dimensional output vector for illustrative purposes, resulting in an output vector 2 x 7. These coordinates serve as an objective measure of the motif's relationship to every other motif in the dataset. The coordinates and the one-hot encoded sequences are fed into respective CNNs. The resulting transformed matrices are concatenated and fed into a series of fully-connected layers, and the network provides a one-hot label vector as output.

kinase-domain sequences, that is, the specific subsequences of kinases that are directly involved in phosphorylation. This distance was 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,

$$d(f_a, f_b) = \frac{\sum_{k_a \in f_a} \sum_{k_b \in f_b} Lev(k_a, k_b)}{|f_a| \cdot |f_b|} \tag{5}$$

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 x b kinase-family dissimilarity matrix. Here, we also refer to this proxy metric for evolutionary dissimilarity between kinase families as the "phylogenetic distance".

Kinase-family dissimilarity vs. 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 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 5, but based on the motifs specific to each kinase family, resulting in an a x b motif-group dissimilarity matrix.

Kinase phylogenetic loss. To leverage evolutionary relationships between kinase families into our predictions, we constructed a mean squared error (MSE) loss function weighted by a kinase phylogenetic metric. Specifically, our weighted MSE loss per minibatch is defined as:

$$PMSE(\hat{y}, y) = \frac{1}{n} \sum_{i=1}^{n} P_i^T (y_i - \hat{y}_i)^2, \tag{6}$$

where n is the size of the mini batch, y_i is the one-hot actual 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

$$P_i = \left[w_{0,i}, ..., w_{|K|,i}\right]^T, \tag{7}$$

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

$$w_{k,i} = \frac{1}{|L_i|} \sum_{j \in L_i} F_{k,j}, \tag{8}$$

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

$$L_i = \{i \in [0, ..., m-1] : y_i = 1\}, \tag{9}$$

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

We note that our phylogenetic weights within the MSE loss performed better than the standard binary cross-entropy (BCE) loss usually employed for classification tasks. We hypothesize that MSE's improved performance has to do with

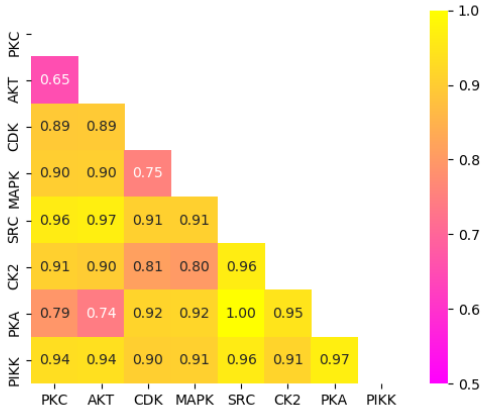

Fig. 3. Heat map matrix depicting pairwise kinase-domain Levenshtein distances. The distances were normalized; the yellow end of the color bar represents far distances, and the pink end of the color bar represents close distances. The distances were normalized, but the original set of distances included distances between kinase families and themselves, i.e. $\mathrm{Lev}(\mathrm{PKA}, \mathrm{PKA})$, which are very small. As a result, all remaining pairwise distances, i.e. $\mathrm{Lev}(K_a, K_b)$ where $a \neq b$ were pushed to distances between 0.65 and 1.00. Note that the distances are unitless.

information lost when applying the BCE loss in comparison to the MSE loss. Consider the BCE loss equation:

$$BCE(\hat{y}, y) = -\sum_i y_i \log(\hat{y}_i). \quad \textbf{(10)}$$

Here, any actual negative label ($y = 0$) will not be integrated into the loss, thus losing information that may aid in model training.

Results

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 the Methods section. 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. While moderate, this correlation between kinase dissimilarity and motif dissimilarity suggests a potential signal in the phylogenetic relationships that could be leveraged to improve predictions.

Using our normalized distances as a proxy for phylogenetic distance (see Methods), the dissimilarity between kinases is displayed as a heatmap in Figure 3. The Akt and PKC family 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 include both motif dissimilarity and kinase relatedness into the predictive model, as achieved through our custom phylogenetic loss function described in Methods, and the effects of which are described later in Results.

Table 2. Precision and recall scores of independent test set predictions, given by kNN performed on the ProtVec and Siamese embeddings.

	Precision		Recall	
Family	ProtVec	Siamese	ProtVec	Siamese
Akt	0.463	0.463	0.500	0.500
CDK	0.933	0.933	0.504	0.504
CK2	0.849	0.911	0.735	0.706
MAPK	0.833	0.846	0.791	0.748
PIKK	0.463	0.964	0.500	0.508
PKA	0.829	0.815	0.574	0.599
PKC	0.744	0.812	0.512	0.562
Src	0.968	0.999	0.737	0.955
average	0.760	**0.843**	0.607	**0.634**

Motif embedding via Siamese network. We also sought to develop a novel learned representation of motifs using a Siamese neural network. Siamese networks were first introduced in the early 1990s as a method to solve signature verification, posed as an image-to-image matching problem (23). Siamese networks perform metric learning by exploiting the dissimilarity between a pair of data points. Training a Siamese 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 model, we opted to use a single bidirectional long short-term memory (LSTM) layer, followed by fully-connected layers for each twin, as we found that the bidirectional LSTM architecture provided superior results when compared to the convolutional neural networks (CNN) alternative (data not shown). We performed k-nearest neighbors 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 an average precision of 0.843 compared to ProtVec's 0.760. Likewise, the Siamese embedding had better recall, with an average recall of 0.634 compared to ProtVec's 0.607.

We performed dimensionality reduction for visualization of the Siamese embeddings using uniform manifold approximation and projection (UMAP) (25). For our UMAP implementation, we used 200 neighbors, a minimum distance of 0.1, and Euclidean distance for our metric. The resulting 2-dimensional UMAP motif embeddings derived from the Siamese network are shown in Figure 4. As can be seen, the motifs phosphorylated by a given kinase family have a distinctive distribution in the embedding space, with some distributions being highly unique and with significant overlap between other families. More specifically, our Siamese embedding shows that motifs phosphorylated by either PKC, PKA, or Akt appear to occupy a similar latent space. Sim-

ilarly, 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 8 kinase families we investigated, with its motifs invariably having a tyrosine (Y) at the eighth position in the 15-amino acid sequence, compared to the other 7 families whose motifs have either a serine (S) or a (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 motif-kinase phosphorylation events instead of motif sequence comparisons. We note that the embedding produced by our Siamese network is quite qualitatively similar to the ProtVec embedding in terms of these kinase-label clusters indicated in the UMAP projections.

Prediction of phosphorylation events. Following training of EMBER with 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 associated CNN layers; this was achieved by applying a dropout rate of 1.00 on the final layer of the coordinate associated CNN. This ablation test allowed us to observe how our novel motif sequence-coordinate model compares to a canonical deep learning model whose input consists of solely one-hot encoded motif sequences (such as in the methods utilized by (17) and (18)). We also compared EMBER trained on the standard BCE loss to EMBER trained on our kinase phylogenetic loss. All predictive models, as described in Table 3, were trained on the same training set and evaluated on the same independent test set.

Comparisons between the predictive capability of the models described here are quantified as area under the ROC curve (AUROC) and area under the precision-recall curve (mean precision), and these metrics are presented for each of the four 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 mean precision. In addition, integration of the phylogenetic loss provides a generally small but consistent additional boost in performance, showing the best overall results out of the three models for AUROC and mean precision. Individual metric curves for each kinase label, produced by EMBER trained via the phylogenetic loss, are shown in 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 pre-

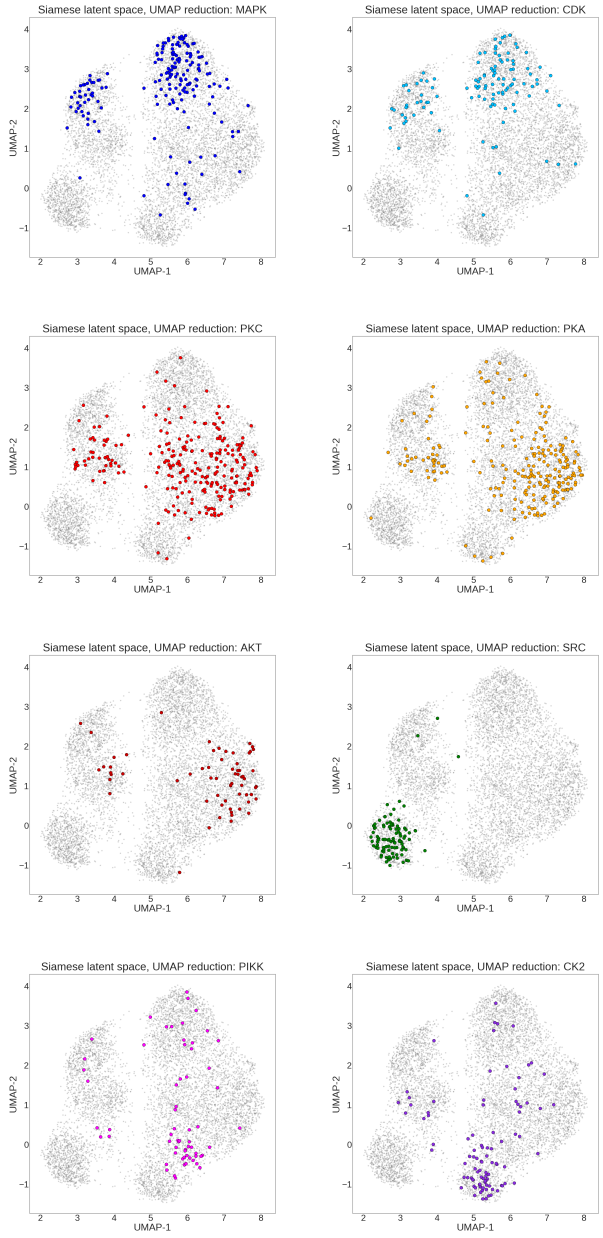

Fig. 4. Siamese embedding of motifs. Each point represents one of the 7535 motifs, and each panel displays kinase-specific phosphorylation patterns, with each colored point corresponding to a motif in the test set phosphorylated by the specified kinase. Highlighted points are slightly enlarged in size to enhance readability.

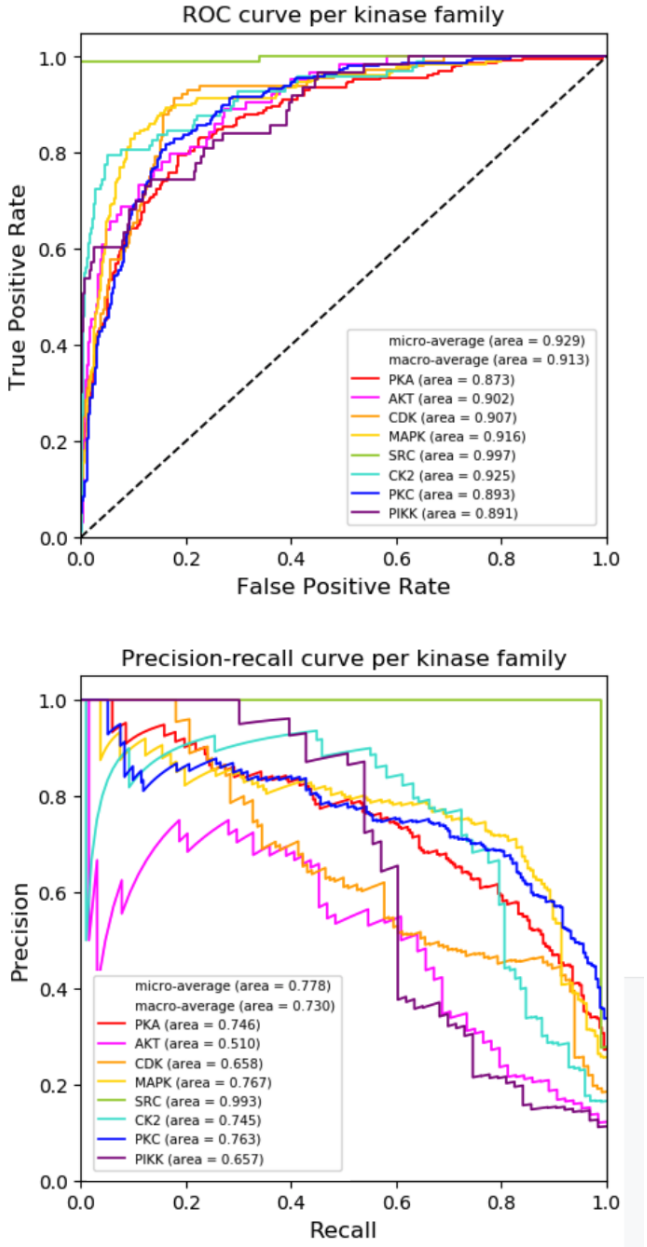

Fig. 5. Receiver operator curves and precision-recall curves per kinase label, resulting from EMBER.

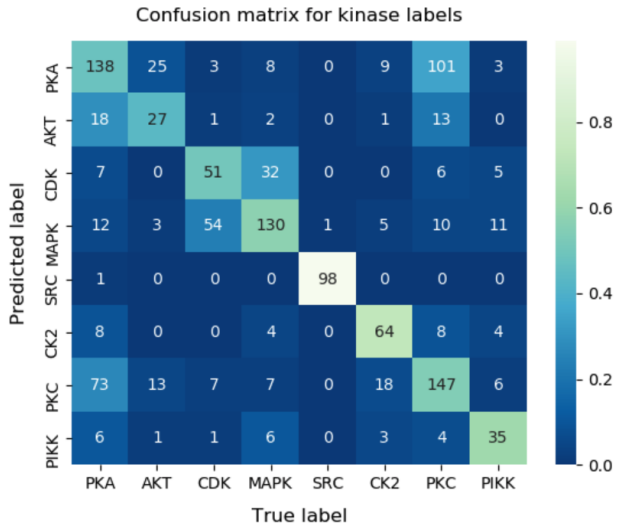

Fig. 6. Confusion matrix for EMBER predictions on the test set. The numbers inside each box represent the raw number of predictions per box. The color scale is based on the ratio of predictions (in the corresponding box) to total predictions, per label. A lighter color corresponds to a larger ratio of predictions to total predictions. Therefore, even though (Src, Src = 98) is the lightest color in the matrix, it is not the highest value in the matrix.

diction equal to or less than 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 (Figure 3). Furthermore, our Siamese 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.

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, for example, by facilitating the development of novel kinase inhibitor-based therapies that disrupt kinase signaling pathways. Here, we have presented a deep learning-based approach that aims to predict which substrates are likely to be phosphorylated by a specific kinase. 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

Table 3. AUROC and mean precision results achieved on the independent test set across classification models. The area under the receiver operating characteristic (AUROC) and area under the precision-recall curve (mean precision) are presented per kinase family for each classification model. From left to right, we include results for the initial kNN performed on the Siamese embedding, the ablated sequence-only CNN, EMBER trained using a BCE loss, and EMBER trained using the kinase phylogenetic MSE loss as described in Methods.

	AUROC				Mean precision			
Family	kNN	Seq-CNN	EMBER (*BCE*)	EMBER (*P-MSE*)	kNN	Seq-CNN	EMBER (*BCE*)	EMBER (*P-MSE*)
Akt	0.886	0.892	0.910	0.902	0.460	0.518	0.509	0.510
CDK	0.876	0.900	0.906	0.907	0.488	0.639	0.642	0.658
CK2	0.914	0.915	0.919	0.925	0.694	0.745	0.747	0.745
MAPK	0.904	0.895	0.907	0.916	0.750	0.694	0.764	0.767
PIKK	0.863	0.878	0.892	0.891	0.598	0.645	0.652	0.657
PKA	0.851	0.866	0.865	0.873	0.696	0.724	0.723	0.746
PKC	0.879	0.886	0.891	0.893	0.721	0.738	0.754	0.763
Src	0.999	0.997	0.997	0.997	0.994	0.993	0.993	0.993
micro-average	0.917	0.886	0.921	**0.929**	0.720	0.652	0.731	**0.778**
macro-average	0.897	0.904	0.911	**0.913**	0.675	0.712	0.723	**0.730**

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 kNN task. This improvement over ProtVec was in spite of the fact that the Siamese network utilized less than 7,000 training sequences of 15 amino acids in length compared to ProtVec's 500,000 sequences of approximately 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 network to be trained on any number of multi-label datasets such acetylation, methylation, and 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. This work was supported by grants through the National Institutes of Health (Grant #s CA177993, CA233811, CA238475, DK116204 and DK109559).

Bibliography

1. Tzong-Yi Lee, Justin Bo-Kai Hsu, Wen-Chi Chang, and Hsien-Da Huang. RegPhos: a system to explore the protein kinase-substrate phosphorylation network in humans. *Nucleic Acids Res.*, 39(Database issue):D777–87, January 2011.
2. G Manning, D B Whyte, R Martinez, T Hunter, and S Sudarsanam. The protein kinase complement of the human genome. *Science*, 298(5600):1912–1934, December 2002.
3. Panayotis Vlastaridis, Pelagia Kyriakidou, Anargyros Chaliotis, Yves Van de Peer, Stephen G Oliver, and Grigoris D Amoutzias. Estimating the total number of phosphoproteins and phosphorylation sites in eukaryotic proteomes. *Gigascience*, 6(2):1–11, February 2017.
4. Leah J Wilson, Adam Linley, Dean E Hammond, Fiona E Hood, Judy M Coulson, David J MacEwan, Sarah J Ross, Joseph R Slupsky, Paul D Smith, Patrick A Eyers, and Ian A Prior. New perspectives, opportunities, and challenges in exploring the human protein kinome. *Cancer Res.*, December 2017.
5. G L Johnson and Razvan Lapadat. Mitogen-Activated protein kinase pathways mediated by ERK , JNK , and p38 protein kinases. *Science*, 298(5600):1911, 2002.
6. Gayathri K Perera, Chrysanthi Ainali, Ekaterina Semenova, Christian Hundhausen, Guillermo Barinaga, Deepika Kassen, Andrew E Williams, Muddassar M Mirza, Mercedesz Balazs, Xiaoting Wang, Robert Sanchez Rodriguez, Andrej Alendar, Jonathan Barker, Sophia Tsoka, Wenjun Ouyang, and Frank O Nestle. Integrative biology approach identifies cytokine targeting strategies for psoriasis. *Sci. Transl. Med.*, 6(223):223ra22, February 2014.
7. Nicole Tegtmeyer, Matthias Neddermann, Carmen Isabell Asche, and Steffen Backert. Subversion of host kinases: a key network in cellular signaling hijacked by helicobacter pylori CagA. *Mol. Microbiol.*, May 2017.
8. Amandine Charras, Pinelopi Arvaniti, Christelle Le Dantec, Marina I Arleevskaya, Kaliopi Zachou, George N Dalekos, Anne Bordon, and Yves Renaudineau. JAK inhibitors suppress innate epigenetic reprogramming: a promise for patients with sjögren's syndrome. *Clin. Rev. Allergy Immunol.*, June 2019.
9. Alessia Alunno, Ivan Padjen, Antonis Fanouriakis, and Dimitrios T Boumpas. 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(8), August 2019.
10. Ya Nan Deng, Joseph A Bellanti, and Song Guo Zheng. Essential kinases and transcriptional regulators and their roles in autoimmunity. *Biomolecules*, 9(4), April 2019.
11. Kyla A L Collins, Timothy J Stuhlmiller, Jon S Zawistowski, Michael P East, Trang T Pham, Claire R Hall, Daniel R Goulet, Samantha M Bevill, Steven P Angus, Sara H Velarde, Noah Sciaky, Tudor I Oprea, Lee M Graves, Gary L Johnson, and Shawn M Gomez. Proteomic analysis defines kinase taxonomies specific for subtypes of breast cancer. *Oncotarget*, 9 (21):15480–15497, March 2018.
12. Elise J Needham, Benjamin L Parker, Timur Burykin, David E James, and Sean J Humphrey. Illuminating the dark phosphoproteome. *Sci. Signal.*, 12(565), January 2019.
13. Wenwen Fan, Xiaoyi Xu, Yi Shen, Huanqing Feng, Ao Li, and Minghui Wang. Prediction of protein kinase-specific phosphorylation sites in hierarchical structure using functional information and random forest. *Amino Acids*, 46(4):1069–1078, April 2014.
14. Shu-Yun Huang, Shao-Ping Shi, Jian-Ding Qiu, and Ming-Chu Liu. Using support vector machines to identify protein phosphorylation sites in viruses. *J. Mol. Graph. Model.*, 56: 84–90, March 2015.
15. Fuyi Li, Chen Li, Tatiana T Marquez-Lago, André Leier, Tatsuya Akutsu, Anthony W Purcell, A Ian Smith, Trevor Lithgow, Roger J Daly, Jiangning Song, and Kuo-Chen Chou. Quokka: a comprehensive tool for rapid and accurate prediction of kinase family-specific phosphorylation sites in the human proteome. *Bioinformatics*, 34(24):4223–4231, December 2018.
16. Yu Xue, Ao Li, Lirong Wang, Huanqing Feng, and Xuebiao Yao. PPSP: prediction of PK-specific phosphorylation site with bayesian decision theory. *BMC Bioinformatics*, 7:163, March 2006.
17. Duolin Wang, Shuai Zeng, Chunhui Xu, Wangren Qiu, Yanchun Liang, Trupti Joshi, and Dong Xu. MusiteDeep: a deep-learning framework for general and kinase-specific phosphorylation site prediction. *Bioinformatics*, 33(24):3909–3916, December 2017.
18. Fenglin Luo, Minghui Wang, Yu Liu, Xing-Ming Zhao, and Ao Li. DeepPhos: prediction of protein phosphorylation sites with deep learning. *Bioinformatics*, 35(16):2766–2773, August 2019.
19. Ehsaneddin Asgari and Mohammad R K Mofrad. Continuous distributed representation of biological sequences for deep proteomics and genomics. *PLoS One*, 10(11):e0141287, November 2015.
20. Peter V Hornbeck, Jon M Kornhauser, Sasha Tkachev, Bin Zhang, Elzbieta Skrzypek,

Beth Murray, Vaughan Latham, and Michael Sullivan. 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): D261–70, January 2012.

21. Jianfei Hu, Hee-Sool Rho, Robert H Newman, Jin Zhang, Heng Zhu, and Jiang Qian. PhosphoNetworks: a database for human phosphorylation networks. *Bioinformatics*, 30(1):141–142, January 2014.
22. Holger Dinkel, Claudia Chica, Allegra Via, Cathryn M Gould, Lars J Jensen, Toby J Gibson, and Francesca Diella. Phospho.ELM: a database of phosphorylation sites–update 2011. *Nucleic Acids Res.*, 39(Database issue):D261–7, January 2011.
23. Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a "siamese" time delay neural network. In J D Cowan, G Tesauro, and J Alspector, editors, *Advances in Neural Information Processing Systems 6*, pages 737–744. Morgan-Kaufmann, 1994.
24. R Hadsell, S Chopra, and Y LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742, June 2006.
25. Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. February 2018.