

Graph Space Embedding

João Pereira^{1,2}, Albert K. Groen¹, Erik S. G. Stroes¹ and Evgeni Levin^{1,2}

¹Amsterdam University Medical Center, The Netherlands

²Horaizon BV, The Netherlands

{j.p.belopereira, e.stroes, a.k.groen, e.levin}@amsterdamumc.nl

Abstract

We propose the *Graph Space Embedding* (GSE), a technique that maps the input into a space where interactions are implicitly encoded, with little computations required. We provide theoretical results on an optimal regime for the GSE, namely a feasibility region for its parameters, and demonstrate the experimental relevance of our findings. Next, we introduce a strategy to gain insight on which interactions are responsible for the certain predictions, paving the way for a far more transparent model. In an empirical evaluation on a real-world clinical cohort containing patients with suspected coronary artery disease, the GSE achieves far better performance than traditional algorithms.

1 Introduction

Learning from interconnected systems can be a particularly difficult task due to the possibly non-linear interaction between the components [Linde *et al.*, 2015; Bereau *et al.*, 2018]. In some cases, these interactions are known and therefore constitute an important source of prior information [Jonshkowski, 2015; Zhou *et al.*, 2018]. Although prior knowledge can be leveraged in a variety of ways [Yu *et al.*, 2010], most of the research involving interactions, is focused on their discovery. One popular approach to deal with feature interactions, is to cast the interaction network as a graph and then use kernel methods based on graph properties, such as walk-lengths or subgraphs [Borgwardt and Kriegel, 2005; Shervashidze *et al.*, 2009; Kriege and Mutzel, 2012] or, more recently, graph deep convolutional methods [Defferrard *et al.*, 2016; Fout *et al.*, 2017; Kipf and Welling, 2017]. In this work however, we focus on the case in which the interactions are feature specific and a universal property of the data instances, which make the pattern search algorithms not suitable for this task. To our knowledge, there is limited research involving this setting, although we suggest many problems can be formulated in the same way (see Figure 1). To address this knowledge gap, we present a novel method: *Graph Space Embedding* (GSE), an approach related to the ‘random-walk’ graph kernel [Gärtner *et al.*, 2003; Kang *et al.*, 2012] with an important difference: it is not limited to the sum of all walks of a given length, but rather compares similar edges

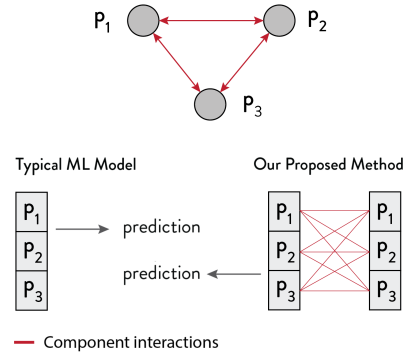


Figure 1: A traditional learning algorithm with no structural information will take the feature values and learn to produce a prediction with complete disregard for their interactions (top graph).

in two different graphs, which results in better expressiveness. Our empirical evaluation demonstrates that GSE leads to an improvement in performance compared to other baseline algorithms when plasma protein measurements and their interactions are used to predict ischaemia in patients with Coronary Artery Disease (CAD) [van Nunen *et al.*, 2015; Zimmermann *et al.*, 2015]. Moreover, the kernel can be computed in $\mathcal{O}(n^2)$, where n is the number of features, and its hyperparameters efficiently optimized via maximization of the kernel matrix variation.

1.1 Main Contributions

1. *Graph Space Embedding* function that efficiently maps input into an “interaction-based” space
2. Novel theoretical result on optimal regime for the GSE, namely feasibility region for its parameters
3. *Even Decent Sampling Algorithm*: a strategy to gain insight on which interactions are responsible for the certain prediction

2 Approach

A remark on notation: we will use bold capital letters for matrices, bold letters for arrays and lower case letters for scalars/functions/1-d variables (ex. \mathbf{X} , \mathbf{x} , x).

2.1 Interaction Graphs

Any network can be represented by a graph $\mathcal{G} = \{V, E\}$, where E is a set of edges, V a set of vertices. Denote by $\mathbf{A}_{|V| \times |V|}$ ($|V|$ is equal to the number of features N) the adjacency matrix, where $\mathbf{A}_{i,j}$ represents the interaction between feature i and j , and whose value is 0 if there is no interaction.

Let $\mathbf{x}_{1 \times N}$ be an array with measurements of features 1 to N for a given point in the data. In order to construct an instance-specific matrix, one can weigh the interaction between each pair of features with a function of their values' product:

$$\mathbf{G}_{\mathbf{x}}(\mathbf{A}) = \varphi(\mathbf{A}) \circ \mathbf{x}^T \mathbf{x}, \quad (1)$$

where $\varphi(\mathbf{A})$ is some function of the network interaction matrix \mathbf{A} , and the operator \circ represents the Hadamard product, i.e. $(\mathbf{A} \circ \mathbf{B})_{i,j} = (\mathbf{A})_{i,j}(\mathbf{B})_{i,j}$.

2.2 Graph Kernel

Unlike the distance in euclidean geometry, which intuitively represents the length of a line between two points, there is no such tangible metric for graphs. Instead, one has to decide what is a reasonable evaluation for the difference between two graphs in the context of the problem.

A popular approach [Gärtner *et al.*, 2003] is to compare random walks on both graphs. The i, j th entry of the order k power of an adjacency matrix $\mathbf{A}_{|V| \times |V|}$: $\mathbf{A}^k = \underbrace{\mathbf{A}\mathbf{A}\dots\mathbf{A}}_{k \text{ times}}$ corresponds to the number of walks of length k from i to j . Any function that maps the data into a feature space \mathcal{H} : $\phi: X \rightarrow \mathcal{H}$, $k(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle$ is a kernel function. Using the original graph kernel formulation, it is possible to define a kernel that will implicitly map the data into a space where the interactions are incorporated:

$$k_n(\mathbf{G}, \mathbf{G}') = \sum_{i,j=1}^n [\gamma]_{i,j} \langle [\mathbf{G}]^i, [\mathbf{G}']^j \rangle_F, \quad (2)$$

where \mathbf{G} and \mathbf{G}' correspond to $\mathbf{G}_{\mathbf{x}}(\mathbf{A})$ and $\mathbf{G}_{\mathbf{x}'}(\mathbf{A})$ (see eq. 1); $\gamma_{i,j}$ is a function that "controls" the mapping $\phi(\cdot)$; and n is the maximum allowed "random walks" length. If γ is decomposed into $\mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, where \mathbf{U} is a matrix whose columns are the eigenvectors of γ , and $\mathbf{\Lambda}$ a diagonal matrix with its eigenvalues at each diagonal entry, then equation 2 can be re-factored into:

$$k_n(\mathbf{G}, \mathbf{G}') = \sum_{k,l=1}^{|V|} \sum_{i=1}^n \phi_{i,k,l}(\mathbf{G}) \phi_{i,k,l}(\mathbf{G}'), \quad (3)$$

where $\phi_{i,k,l}(\mathbf{G}) = \sum_{j=1}^n [\sqrt{\mathbf{\Lambda}\mathbf{U}^T}]_{i,j} \mathbf{G}^j$. Consequently, different forms of the function γ can be chosen, with different interpretations. For the case where $\gamma_{i,j} = \theta^i \theta^j$, which yields:

$$\begin{aligned} k_n(\mathbf{G}, \mathbf{G}') &= \left\langle \sum_{i=1}^n \theta^i [\mathbf{G}]^i, \sum_{j=1}^n \theta^j [\mathbf{G}']^j \right\rangle_F \\ &= \left\langle \sum_{i=1}^n \theta^i [\mathbf{G}]^i, \sum_{i=1}^n \theta^i [\mathbf{G}']^i \right\rangle_F, \end{aligned} \quad (4)$$

the kernel entry can be interpreted as an inner product in a space where there is a feature for every node pair $\{k, l\}$,

which represents the weighted sum of paths of length 1 to n from k to l ($\phi_{k,l} = \sum_{i=1}^n \theta^i \mathbf{G}_{k,l}^i$) [Tsivtsivadze *et al.*, 2011]. The kernel can then be used with a method that employs the kernel trick, such as support vector machines, kernel PCA or kernel clustering. Another interesting case is when we consider the weighted sum of paths of length 1 to ∞ . This can be calculated using:

$$k_{\infty}(\mathbf{G}, \mathbf{G}') = \langle e^{\beta \mathbf{G}}, e^{\beta \mathbf{G}'} \rangle_F, \quad (5)$$

since $e^{\beta \mathbf{G}} = \lim_{n \rightarrow +\infty} \sum_{i=0}^n \frac{\beta^i}{i!} \mathbf{G}^i$, where β is a parameter.

2.3 Graph Space Embedding

Since we are dealing with a universal interaction matrix for every data point and the interactions are feature specific, it makes sense to compare the same set of edges for every pair of points. As a consequence, we can also avoid solving time-consuming graph structure problems. With these two points in mind, we combined the previous graph kernel methods and the radial basis function (RBF) to develop a new kernel which we will henceforth refer to as Graph Space Embedding (GSE). The radial basis function is defined as:

$$k(\mathbf{x}, \mathbf{y}) = e^{-\frac{\|\mathbf{x}-\mathbf{y}\|^2}{\sigma^2}} = c e^{\frac{2\langle \mathbf{x}, \mathbf{y} \rangle}{\sigma^2}}, \quad (6)$$

where $c = e^{-\frac{\|\mathbf{x}\|^2}{\sigma^2}} e^{-\frac{\|\mathbf{y}\|^2}{\sigma^2}}$. The GSE uses the distance $\langle \sqrt{\gamma}[\mathbf{G}], \sqrt{\gamma}[\mathbf{G}'] \rangle_F$ in the radial basis function:

$$k(\mathbf{G}, \mathbf{G}') = c e^{\frac{2\langle \mathbf{x}, \mathbf{y} \rangle}{\sigma^2}} = c \underbrace{\sum_{n=0}^{\infty} \frac{(2 \langle \sqrt{\gamma} \mathbf{G}, \sqrt{\gamma} \mathbf{G}' \rangle_F)^n}{\sigma^{2n} n!}}_{r_w} \quad (7)$$

If we then take the upper term of the fraction in r_w to be $\left[2 \sum_{i=0}^{|E|} \gamma \mathbf{G}_i \mathbf{G}'_i \right]^n$, we can use the multinomial theorem to expand each term of the exponential power series, and the expression for the kernel then becomes:

$$k(\mathbf{G}, \mathbf{G}') = c \underbrace{\sum_{n=0}^{\infty} \left(\frac{2}{\nu} \right)^n}_{\lambda} \underbrace{\sum_{\alpha^n(\cdot)} \frac{\prod_{i=1}^{|E|} [\mathbf{G}_i \mathbf{G}'_i]^{\alpha_i}}{\prod_{i=1}^{|E|} \Gamma(\alpha_i + 1)}}_{r_e}, \quad (8)$$

where Γ is the gamma function, $\mathbf{G}_i \in E$ is the value of edge i in \mathbf{G} and $\nu = \frac{\sigma^2}{\gamma}$. Here, $\alpha^n(\cdot)$ represents a combination of $|E|$ integers: $(\alpha_1, \alpha_2, \dots, \alpha_{|E|})$, with $\sum_i \alpha_i = n$, and the sum in r_e is taken over all possible combinations of $\alpha^n(\cdot)$. For instance, for $n = 3$ in a graph with $|E| = 5$, possible examples of $\alpha^3(\cdot)$ include $(0, 1, 1, 1, 0)$ or $(0, 2, 1, 0, 0)$ (see Figure 2). We begin by noting that since the sum in r_e is taken over all combinations $(l, k) \in V \times V$ of size n , the GSE then represents a mapping from the input space to a space where all combinations of $n = 0 \rightarrow \infty$ edges are compared between \mathbf{G} and \mathbf{G}' , walks or otherwise (see fig 2). Notice that this is in contrast with the kernel of equation 5, where the comparison is between a sum of all possible walks of length $n = 0 \rightarrow \infty$ from one node to another in the two graphs.




Figure 2: The GSE kernel implicitly compares all edge combinations between \mathbf{G} and \mathbf{G}' . In this hypothetical graph, we show a sample of four α combinations for $n = 3$. We denote by $r.e(\alpha(i))$ the value inside the sum $r.e$ (see eq. 8) corresponding to the combination $\alpha(i)$. Note that while $\alpha(1)$ is a graph walk and $\alpha(2)$ is not, $r.e(\alpha(1)) = r.e(\alpha(2))$. However, due to the repetitions in $\alpha(3)$ and $\alpha(4)$, their value is shrunk in relation to the others. The higher the number of repetitions, the more the value shrinks.

The GSE also allows repeated edges. However, if the data is normalized so that $\mu(\mathbf{G}_i) \simeq 0, \sigma(\mathbf{G}_i) \simeq 1$, then both the power in the numerator and the denominator of $r.e$ will effectively dampen most combinations with repeated edges, with a higher dampening factor for higher number of repetitions and/or combinations. Even for outlier values, the gamma function will quickly dominate the numerator of $r.e$. The λ factor serves the purpose of shrinking the combinations with higher number of edges for $\nu > 2$. Finally, σ^2 now serves a dual purpose: the usual one in RBF to control the influence of points in relation to their distance (see equation 6), while at the same time controlling how much combinations of increasing order are penalized.

2.4 ν Feasibility Region

As discussed in the above section, the hyperparameter ν controls the shrinking of the contribution of higher order edge combinations. Intuitively, not all values of ν will yield a proper kernel matrix since too large of a value will leave out too many edge combinations while one too small will saturate the kernel values. This motivates the search for a ν value feasible operation region, where the kernel incorporates the necessary information for separability. Informally speaking, the kernel entry $k(\mathbf{G}, \mathbf{G}')$ measures the similarity of \mathbf{G} and \mathbf{G}' . In case too few/many edge combinations are considered, the variation of the kernel values will be equal to 1. Therefore, we use the variation of the kernel matrix $\sigma^2(\mathbf{K})$ as a proxy to detect if ν is within acceptable bounds. We shall refer to the ability of the kernel to map the points in the data into separable images $\phi(\mathbf{x})$ as kernel expressiveness.

To determine this region analytically, we find the ν_{max} that yields the largest kernel variation, and then use the loss function around this value to determine in which direction the value ν should take for minimal loss.

Lemma 2.1. $\max_{\nu} \sigma^2(\mathbf{K}(\nu))$ can be numerically estimated and is guaranteed to converge with a learning rate $\alpha \leq$

$\frac{D}{2(D-1)d_{max}}$, where D is the total number of inter graph combinations and d_{max} is the largest combination distance.

Proof. The analytical expression for the variance is:

$$\sigma^2(\mathbf{K}(\nu)) = E[\mathbf{K}(\nu)^2] - \underbrace{E[\mathbf{K}(\nu)]^2}_b = \left(\frac{D-1}{D}\right) \sum_{d=1}^D e^{-2\nu d} - \frac{1}{D^2} \sum_{i \neq j}^{D^2-D} 2e^{-\nu(d_i+d_j)}, \quad (9)$$

where we used the binomial theorem to expand b , and $d = \|\mathbf{G} - \mathbf{G}'\|^2$. To guarantee the convergence of numerical methods the function derivative must be Lipschitz continuous:

$$\frac{\|\mathbf{K}'(\nu) - \mathbf{K}'(\nu')\|}{\|\nu - \nu'\|} \leq L(\mathbf{K}') : \forall \nu, \nu', \quad (10)$$

by overloading the notation: $\mathbf{K}'(\nu) = \frac{\partial \sigma^2(\mathbf{K}(\nu))}{\partial \nu}$ to simplify the expression. The left side of equation 10 becomes:

$$\frac{\|\top - \Lambda\|}{\|\nu - \nu'\|}, \quad \Lambda = 2 \left(\frac{D-1}{D^2}\right) \left[\sum_{d=1}^D d(e^{-2\nu d} - e^{-2\nu' d}) \right], \quad \top = \frac{2}{D^2} \sum_{i \neq j}^{D^2-D} (d_i + d_j) \left(e^{-\nu(d_i+d_j)} - e^{-\nu'(d_i+d_j)} \right). \quad (11)$$

Since $0 \leq e^{-\beta} \leq 1 : \forall \beta \in \mathbb{R}$, then:

$$\|\top - \Lambda\| \leq 2 \left(\frac{D-1}{D}\right) d_{max}. \quad (12)$$

When $\epsilon = \nu - \nu' \rightarrow 0$:

$$e^{-c\nu} - e^{-c\nu'} = \frac{e^{c\nu'} - e^{c\nu}}{\underbrace{e^{c(\nu+\nu')}}_{\delta}} \rightarrow 0, : \nu, \nu' > 0, \quad (13)$$

and δ tends much faster to 0 than ϵ , since the denominator of δ is the exponential of the sum of ν and ν' . Thus, the function $k'(\nu)$ is Lipschitz continuous with constant equal to: $L(\mathbf{K}'(\nu)) = 2 \left(\frac{D-1}{D}\right) d_{max}$. \square

We shall later demonstrate empirically that $\nu^* = \max_{\nu} \sigma^2(\mathbf{K}(\nu))$ improves the class separability for our dataset.

2.5 Comparison with Standard Graph Kernels

The original formulation of the graph kernel by Gartner et. al (see eq. 2), multiplies sums of random walks of length i from one edge to another ($k \rightarrow l$) by sums of random walks $k \rightarrow l$ from the other graph being compared of a length not necessarily equal to i :

$$k_n(\mathbf{G}, \mathbf{G}') = \sum_{i,j=1}^n [\gamma]_{i,j} \left\langle [\mathbf{G}]_{kl}^i, [\mathbf{G}']_{kl}^j \right\rangle_F = \sum_{k,l=1}^{|V|} \sum_{i=1}^n [\mathbf{G}]_{kl}^i \sum_{j=1}^n [\gamma]_{i,j} [\mathbf{G}']_{kl}^j. \quad (14)$$

The infinite length random walk formulation (see eq. 5) behaves in a similar way. Our method though, always compares the same set of edges in the two graphs.

Another important difference is the complexity of our method versus the random-walk graph kernel. For an $m \times m$ kernel and $n \times n$ graph, the worst-case complexity for a length k' random walk kernel is $\mathcal{O}(m^2 k' n^4)$ and $\mathcal{O}(m^2 k' n^2)$ for dense and sparse graphs, respectively [Vishwanathan *et al.*, 2010]. The GSE, on the other hand, is always $\mathcal{O}(m^2 n^2)$ since the heaviest operation is the Frobenius inner product in order to compute the distance between \mathbf{G} and \mathbf{G}' . Moreover, once this distance is computed, evaluating the kernel for different values of ν is $\mathcal{O}(1)$, which combined with the fact that the variance of this kernel is Lipschitz continuous, allows for efficient searching of optimal hyperparameters (see section 2.4).

2.6 Interpretability

How could we better understand what the GSE is doing, when it maps points into an infinite-dimensional space? A successful recent development in explaining black-box models is that of Local Interpretable Model-agnostic Explanations (LIME) [Ribeiro *et al.*, 2016], where a model is interpreted locally by making slight perturbations in the input and building an interpretable model around the new predictions. We too shall monitor our model's response to changes in the input, but instead of making random perturbations, we will perturb the input in the direction of maximum output change.

Given an instance from the dataset $\mathbf{x}_{1 \times N}$, where N is the number of features, and the function that will incorporate the feature connection network $\mathbf{G}_x(\mathbf{A})$ (e.g. $\mathbf{G}_x(\mathbf{A}) = \mathbf{A} \circ \mathbf{x}^\top \mathbf{x}$), we will find the direction to which the model is the most sensitive (positive and negative). Unlike optimization, where the goal is to converge as fast as possible, here we are interested in the intermediate steps of the descent. This is because we shall use the set $\mathbf{G} = \{\mathbf{G}_{\mathbf{x}_1}, \mathbf{G}_{\mathbf{x}_2}, \dots, \mathbf{G}_{\mathbf{x}_M}\}$ and the black-box model's predictions $\mathbf{f} = \{f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_M)\}$ to fit our interpretable model $h(\mathbf{G}) \in \mathcal{H}$ (where \mathbf{x}_i is a variation of the original sample \mathbf{x}_0 , and \mathcal{H} represents the space of all possible interpretable functions h). This way, we will indirectly unveil the interactions that our model is most sensitive to, and show how these impact the predictions. To penalize complex models over simpler ones, we will introduce a function $\Omega(h)$ that measures model complexity. To scale the model complexity term appropriately, we can find a scalar θ so that the expected value of $\Omega(h)$ is equal to a fraction ε of the expected value of the loss:

$$\mathbf{E}[\theta \Omega(h)] = \varepsilon \mathbf{E}[\mathcal{L}] \leftrightarrow \theta = \frac{\varepsilon \mathbf{E}[\mathcal{L}]}{\mathbf{E}[\Omega(h)]}. \quad (15)$$

Lastly, for highly non-linear models, the larger the input space the more complex the output explanations are likely to be, so we will weigh the sample deviations the same as the original sample \mathbf{x}_0 using the model's own similarity measure $k(\mathbf{G}_{\mathbf{x}_i}, \mathbf{G}_{\mathbf{x}_0})$. Putting it all together:

$$\xi(\mathbf{x}_0) = \min_{h \in \mathcal{H}} \mathcal{L}(h, f, k(\mathbf{G}_{\mathbf{x}_i}, \mathbf{G}_{\mathbf{x}_0})) + \theta \Omega(h). \quad (16)$$

where $\mathcal{L}(h, f, k(\mathbf{G}_{\mathbf{x}_i}, \mathbf{G}_{\mathbf{x}_0}))$ is the loss of h when using $\mathbf{G}_{\mathbf{x}_i}$ to predict the black-box model output $f(\mathbf{x}_i)$, weighted by the kernel distance to the original sample $k(\mathbf{G}_{\mathbf{x}_i}, \mathbf{G}_{\mathbf{x}_0})$.

Even Descent Sampling Method

In order to adequately cover the most sensitive regions, we need to take steps with equidistant output values. Thus, we developed a novel adaptive method to sample more in steeper regions and less in flatter ones. The intuition is that we would like to approximate the function values in unexplored regions, so that we choose an appropriate sampling step while considering the uncertainty of the approximation. Due to the stochastic nature of the method, it is able to escape local extremes. Consider the value of function f at a point \mathbf{x}_0 and its first order Taylor approximation at an arbitrary point \mathbf{x} :

$$f(\mathbf{x}) \approx \hat{f}(\mathbf{x}) = f(\mathbf{x}_0) + \nabla_{\mathbf{x}} f(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0). \quad (17)$$

The larger the difference $\delta = \mathbf{x} - \mathbf{x}_0$, the less likely it is that the approximation error $f(\mathbf{x}) - \hat{f}(\mathbf{x})$ is small. Assume we would like to model the random variable F , which takes the value of 1 if the approximation error is small ($\delta = |f(\mathbf{x}) - \hat{f}(\mathbf{x})| \approx 0$), and 0 otherwise. We will model the probability density function of F as being:

$$p_F(f = 1|\delta) = \lambda e^{-\lambda \delta}. \quad (18)$$

Consider also the random variable T which takes the value of 1 if the absolute difference in the output for a point \mathbf{x} exceeds an arbitrary threshold ($|f(\mathbf{x}) - f(\mathbf{x}_0)| > \tau$), and 0 otherwise. Assume there is zero probability this event occurs for sufficiently small steps: $\delta < a(\tau)$, for some value $a(\tau)$. Let us further assume that our confidence that $|f(\mathbf{x}) - f(\mathbf{x}_0)| > \tau$ increases linearly after the value $\delta = a(\tau)$, until the maximum confidence level is reached at $\delta = b$. After some value $\delta = c$, we decide not to make any further assumptions about this event, so we attribute zero probability from that point on. This can be modeled as:

$$p_T(t = 1|\delta) = \begin{cases} \frac{2}{v} \frac{\delta - a(\tau)}{u}, & a(\tau) < \delta \leq b \\ \frac{2}{v}, & b < \delta \leq c \\ 0, & \text{otherwise} \end{cases}, \quad (19)$$

where $v = 2c - a(\tau) - b$, $u = b - a(\tau)$ and $T = 1$, if $|f(\mathbf{x}) - f(\mathbf{x}_0)| > \tau$ and 0 otherwise. The distribution of interest is then $p_S = p(f = 1 \cap t = 1|\delta)$. To simplify the calculations, we impose the uncertainty about our approximation (expressed by F) and the likelihood of a sufficiently large output difference (expressed by T) to be independent given δ : $p(f = 1 \cap t = 1|\delta) = p(t = 1|\delta)p(f = 1|\delta)$, and since the goal is to sample steps from this distribution, we will divide it by the normalization constant: $Z = p(f = 1 \cap t = 1) = \int_{-\infty}^{+\infty} p(f = 1 \cap t = 1|\delta) d\delta$. See Figure 3 for an illustration of the method.

There are a couple of properties that can be manipulated for a successful sampling of the output space:

Controlled Termination

To force the algorithm to terminate after a minimum number of samples M_{min} have been sampled, one can decrease the




Figure 3: Illustration of the even descent sampling. $\hat{f}(x)$ approximates the function $f(x)$ and an estimation of how much $\delta = |x - x_0|$ is required to achieve $|f(x) - f(x_0)| \leq \tau$, is computed. Then a sample of x is drawn according to $p_S = p(f = 1 \cap t = 1|\delta)$

value of $a(\tau)$ with each iteration so that it becomes increasingly more likely that a value of δ will be picked such that $|f(\mathbf{x}) - f(\mathbf{x}_0)| < \tau$, terminating the routine. For this purpose, one can compute the estimated threshold value τ_0 that will keep the routine running.

$$|\hat{f}(\mathbf{x}) - f(\mathbf{x}_0)| \geq \tau \Leftrightarrow \sum_{i=1}^N \nabla_{\mathbf{x}} f(\mathbf{x}_0)[i] \delta[i] \geq \tau, \quad (20)$$

where N is the number of features. This is an underdetermined equation, but one possible trivial solution is to set:

$$\delta[i] \geq \frac{\tau}{N' \nabla_{\mathbf{x}} f(\mathbf{x}_0)[i]} \equiv \tau_0, \quad (21)$$

where N' is the number of non-zero gradient values, then let $a(\tau)$ decay with time so that it will reach this limit value after M_{min} iterations:

$$a(\tau)_i = \tau_0 \left(1 + \frac{\theta_a (M_{min} - i)}{M_{min}} \right). \quad (22)$$

Escaping Local Extrema

To make it more likely to escape local extrema, one possibility is to set the cut-off value c larger when the norm of τ_0 (eq. 21) is larger than its expected value, and smaller otherwise:

$$c = b \left(c_l + \frac{\mathbf{E}[\|\tau_0\|_2] - \|\tau_0\|_2}{\mathbf{E}[\|\tau_0\|_2] + \|\tau_0\|_2} \right), \quad c_l \in]2, +\infty[. \quad (23)$$

This formulation allows jumping out from zones where the gradient is locally small, while taking smaller steps where the gradient is larger than expected.

Termination When Too Far from Original Sample

Since we are trying to explain the model locally, the sampling should terminate when the algorithm is exploring too far from the original sample. For that purpose, one can set λ to increase with increasing distance d to the original sample, pushing the probability density towards the left: $\lambda(d) = e^{-\frac{d}{\sigma^2}}$.

Putting all of the above design considerations together, you can find the complete routine in algorithm 1.

Algorithm 1 Even Descent algorithm

Input: $f, \mathbf{x}_0, \mathbf{A}$

Parameter: $\tau, \lambda, \theta_a, b, c_l, M_{min}$

Output: \mathbf{X}', \mathbf{f}

```

1:  $i \leftarrow 0, f_i \leftarrow f(\mathbf{x}_0), \mathbf{f} \leftarrow [f_i], \text{converged} \leftarrow \text{False}$ 
2:  $\mathbf{E}[\|\tau_0\|] = 0, \mathbf{X}' \leftarrow [\mathbf{x}_0]$ 
3: while converged  $\neq$  True do
4:    $i \leftarrow i + 1$ 
5:    $\nabla f \leftarrow \text{ComputePartialDers}(\mathbf{x}_0, \mathbf{A}, f)$ 
6:    $\tau_0 \leftarrow \tau / |N' * \nabla f|$ 
7:    $a, b, c \leftarrow \text{UpdatepS}(i, \theta_a, M_{min}, \mathbf{E}[\|\tau_0\|], c_l)$ 
8:    $\mathbf{E}[\|\tau_0\|] \leftarrow (\mathbf{E}[\|\tau_0\|](i - 1) + \|\tau_0\|) / i$ 
9:    $\delta \leftarrow \text{EvenSample}(\lambda, a, b, c)$ 
10:   $\mathbf{x}_i \leftarrow \mathbf{x}_i \pm \delta * \nabla f$ 
11:  Append( $\mathbf{f}, f(\mathbf{x}_i)$ ), Append( $\mathbf{X}', \mathbf{x}_i$ )
12:  if  $|f_i - f_{i-1}| < \tau$  then
13:    converged  $\leftarrow$  True
14:  end if
15: end while
16: return  $\mathbf{X}', \mathbf{f}$ 

```

3 Experiments

3.1 Materials

For all our analysis, we used plasma protein levels of patients with suspected coronary artery disease who were diagnosed for the presence of ischaemia [Bom *et al.*, 2018]. A total of 332 protein levels were measured using proximity extension arrays [Assarsson *et al.*, 2014], and of the 196 patients, 108 were diagnosed with ischaemia. The protein-protein interactions data is available for download at StringDB [Jensen *et al.*, 2009]. We implemented the GSE and the random walk kernel in python and used sci-kit learn implementation [Pedregosa *et al.*, 2011] for the other algorithms in the comparison.

3.2 Ischaemia Classification Performance

We benchmarked the GSE performance and running time when predicting ischaemia against the random-walk graph kernel, RBF, and random forests. Additionally, in order to test the hypothesis that the protein-interaction information is improving the analysis, we also tested GSE using a constant matrix full of ones as the interaction matrix. For this benchmark, we performed a 10-cycle stratified shuffle cross-validation split on the normalized protein data and recorded the average ROC area under the curve (AUC). To speed up the analysis, we used a training set of 90 pre-selected proteins using univariate feature selection with the F-statistic [Hira and Gillies, 2015]. The results are shown in table 1. The GSE outperformed all the other compared methods, and the fact that the GSE with a constant matrix (GSE*) had a lower performance increases our confidence that the prior interaction knowledge is beneficial for the analysis. The GSE is also considerably faster than the Random-Walk kernel, as expected. To test how both scale increasing feature size, we compared the running time of both for different pre-selected numbers of proteins. The results are depicted in Figure 4.

Method	AUC std	AUC avg	Run time avg(s)
GSE	0.055890	0.814141	7.63
RWGK	0.051704	0.808838	1720
RF	0.066036	0.764141	17.99
GSE*	0.082309	0.787879	6.59
RBF	0.095247	0.779293	1.16

Table 1: The GSE benchmark against random-walk graph kernel (RWGK), random forests (RF), the GSE with constant interaction matrix (GSE*), and radial basis function (RBF). For all kernels, SVM was used as the learning algorithm.

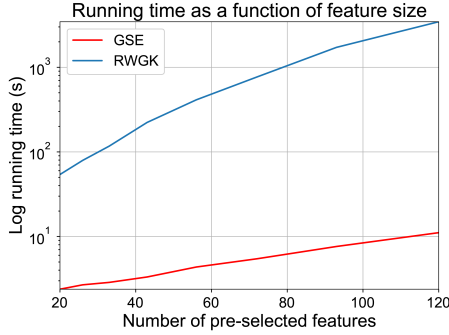


Figure 4: Average running time of the GSE and the Random-walk graph kernel (RWGK), per number of pre-selected features

3.3 Performance for Different ν Values

Recall from section 2.4 that a feasible operating region for the ν values in the GSE kernel was analytically determined. We wanted to investigate how the loss function performs within this region, and whether it is possible to draw conclusions regarding the GSE kernel behaviour with respect to the interactions. To test this, the $\nu^* = \max_{\nu} \sigma^2[k(\nu)]$ was found using a gradient descent (ADAM [Kingma and Ba, 2015]) on the training set over 20 stratified shuffle splits (same pre-processing as in 3.2). We then measured the ROC AUC on

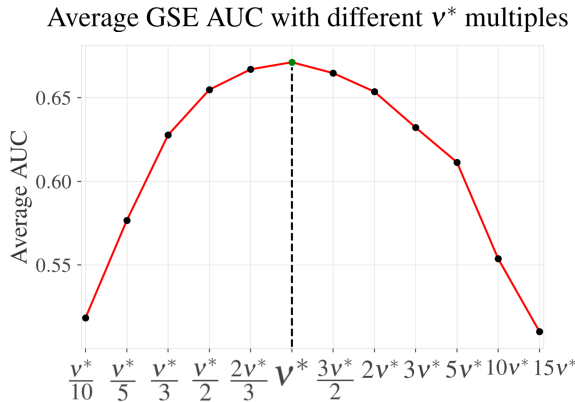


Figure 5: Average ROC AUC on validation set using GSE with different ν values over 20 stratified shuffle splits. Horizontal axis - Multiples of $\max_{\nu} \sigma^2[k(\nu)]$ here denoted by ν^* . The AUC as function of the ν values looks convex and peaks exactly at ν^*

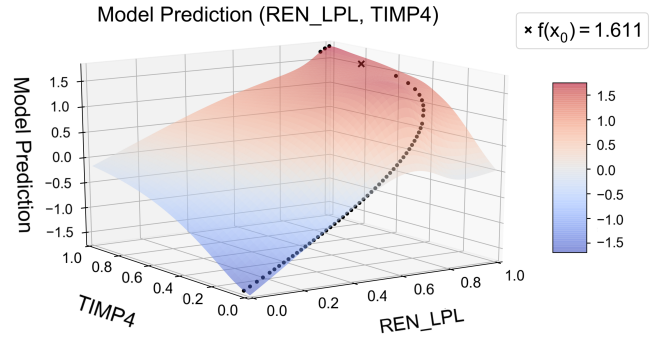


Figure 6: Even Descent Sampling for a random patient in our dataset. This analysis reveals our model "predicts" this patient could be treated by lowering protein "TIMP4" and the interaction between "REN" and "LPL".

the validation set using 12 multiples of ν^* . The results can be seen in Figure 5. It is quite interesting that our proxy for measuring kernel expressiveness turns out to be a convex function peaking at ν^* .

3.4 Interpretability Test

To test how interpretable our model's predictions are, first we trained the model on a random subset of our data and used the trained model to predict the rest of the data. Then we employed the method described in section 2.6 on a random patient in the test set, using decision trees as the interpretable models $h(\mathbf{G}) \in \mathcal{H}$, and a linear weighted combination of *max depth* and *min samples per split* as the complexity penalization term $\Omega(h)$. We then picked the two most important features and made a 3d plot using an interpolation of the prediction space. The result is depicted in Figure 6.

The Even Descent Sampling tests instances which are approximately equidistant in the output values. For this patient, our model 'predicts' its ischaemia risk could be mitigated by lowering protein TIMP metallopeptidase inhibitor 4 ("TIMP4") and the interaction between lipoprotein lipase ("LPL") and renin ("REN").

4 Conclusions

In this paper, we address the problem of analyzing interconnected systems and leveraging the often-known information about how the components interact. To tackle this task, we developed the *Graph Space Embedding* algorithm and compared it to other established methods using a dataset of proteins and their interactions from a clinical cohort to predict ischaemia. The GSE results outperformed the other algorithms in running time and average AUC. Moreover, we presented an optimal regime for the GSE in terms of a feasibility region for its parameters, which vastly decreases the optimization time. Finally, we developed a new technique for interpreting black-box models' decisions, thus making it possible to inspect which features and/or interactions are the most relevant.

References

- [Assarsson *et al.*, 2014] Erika Assarsson, Martin Lundberg, et al. Homogenous 96-plex pea immunoassay exhibiting high sensitivity, specificity, and excellent scalability. *Plos one*, 9(4): e95192, 2014.
- [Bereau *et al.*, 2018] Tristan Bereau, Robert A. DiStasio Jr., Alexandre Tkatchenko, and O. Anatole von Lilienfeld. Non-covalent interactions across organic and biological subsets of chemical space: Physics-based potentials parametrized from machine learning. *The Journal of Chemical Physics*, 148, 2018.
- [Bom *et al.*, 2018] Michiel J. Bom, Evgeni Levin, Paul Knaapen, et al. Predictive value of targeted proteomics for coronary plaque morphology in patients with suspected coronary artery disease. *EBioMedicine.*, 2018.
- [Borgwardt and Kriegel, 2005] Karsten M. Borgwardt and Hans-Peter Kriegel. Shortest-path kernels on graphs. In *In Proceedings of the 5th International Conference on Data Mining*, page 74–81, 2005.
- [Defferrard *et al.*, 2016] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *In Advances in Neural Information Processing Systems.*, page 3844–3852, 2016.
- [Fout *et al.*, 2017] Alex Fout, Jonathon Byrd, Basir Shariat, and Asa Ben-Hur. Protein interface prediction using graph convolutional networks. In *In Advances in Neural Information Processing Systems*, page 6533–6542, 2017.
- [Gärtner *et al.*, 2003] Thomas Gärtner, Peter Flach, and Stefan Wrobel. On graph kernels: Hardness results and efficient alternatives. In *Computational Learning Theory and Kernel Machines, 16th Annual Conference on Computational Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003*, volume 129-143(3), pages 129–143, 2003.
- [Hira and Gillies, 2015] Zena M. Hira and Duncan F. Gillies. A review of feature selection and feature extraction methods applied on microarray data. *Adv Bioinformatics*, 2015.
- [Jensen *et al.*, 2009] Lars J. Jensen, Michael Kuhn, et al. String 8—a global view on proteins and their functional interactions in 630 organisms. *Nucleic Acids Res.*, 37(Database issue):D412-6, 2009.
- [Jonschkowski, 2015] Rico Jonschkowski. Learning state representations with robotic priors. *Autonomous Robots*, 39:407–428, 2015.
- [Kang *et al.*, 2012] U Kang, Hanghang Tong, and Jimeng Sun. Fast random walk graph kernel. In *Proceedings of the 2012 SIAM International Conference on Data Mining*, pages 828–838, 2012.
- [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference for Learning Representations*, 2015.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *In Proceedings of the 6th International Conference on Learning Representations*, 2017.
- [Kriege and Mutzel, 2012] Nils Kriege and Petra Mutzel. Subgraph matching kernels for attributed graphs. In *In Proceedings of the 29th International Conference on Machine Learning*, page 291–298, 2012.
- [Linde *et al.*, 2015] Jörg Linde, Sylvie Schulze, Sebastian G. Henkel, and Reinhard Guthke. Data- and knowledge-based modeling of gene regulatory networks: an update. *EXCLI J.*, 14:346–378, 2015.
- [Pedregosa *et al.*, 2011] Fabian Pedregosa, Gael Varoquaux, et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [Ribeiro *et al.*, 2016] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “why should i trust you?” explaining the predictions of any classifier. In *In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’16).*, 2016.
- [Shervashidze *et al.*, 2009] Nino Shervashidze, S.V.N. Vishwanathan, Tobias H. Petri, Kurt Mehlhorn, and Karsten M. Borgwardt. Efficient graphlet kernels for large graph comparison. In *In Proceedings of the International Conference on Artificial Intelligence and Statistics*, page 488–495, 2009.
- [Tsvitvivadze *et al.*, 2011] Evgeni Tsvitvivadze, Josef Urban, Herman Geuvers, and Tom Heskes. Semantic graph kernels for automated reasoning. In *Proceedings of the Eleventh SIAM International Conference on Data Mining, SDM*, 2011.
- [van Nunen *et al.*, 2015] Lokien X van Nunen, Frederik M Zimmermann, Pim A L Tonino, Emanuele Barbato, Andreas Baumbach, Thomas Engstrøm, et al. Fractional flow reserve versus angiography for guidance of pci in patients with multivessel coronary artery disease (fame): 5-year follow-up of a randomised controlled trial. *Lancet*, 386(10006):1853–1860, 2015.
- [Vishwanathan *et al.*, 2010] S.V. N. Vishwanathan, Nicol N. Schraudolph, Risi Kondor, and Karsten M. Borgwardt. Graph kernels. *Journal of Machine Learning Research*, pages 1201–1242, 2010.
- [Yu *et al.*, 2010] Ting Yu, Simeon Simoff, and Tony Jan. Vqsvm: A case study for incorporating prior domain knowledge into inductive machine learning. *Neurocomputing*, 13-15:2614–2623, 2010.
- [Zhou *et al.*, 2018] Huiwei Zhou, Zhuang Liu, Shixian Ning, Yunlong Yang, Chengkun Lang, Yingyu Lin, and Kun Ma. Leveraging prior knowledge for protein–protein interaction extraction with memory network. *Database*, 2018.
- [Zimmermann *et al.*, 2015] Frederik M. Zimmermann, Angela Ferrara, Nils P. Johnson, Lokien X. van Nunen, Javier Escaned, Per Albertsson, et al. Deferral vs. performance of percutaneous coronary intervention of functionally non-significant coronary stenosis: 15-year follow-up of the deferral trial. *Eur Heart J*, 36:3182–3188, 2015.