

The Path Kernel: A Novel Kernel for Sequential Data

Andrea Baisero, Florian T. Pokorny, Danica Kragic and Carl Henrik Ek

Abstract We define a novel kernel function for finite sequences of arbitrary length which we call the path kernel. We evaluate this kernel in a classification scenario using synthetic data sequences and show that our kernel can outperform state of the art sequential similarity measures. Furthermore, we find that, in our experiments, a clustering of data based on the path kernel results in much improved interpretability of such clusters compared to alternative approaches such as dynamic time warping or the global alignment kernel.

Keywords Kernels · Sequences

1 Introduction

Machine learning methods have had an enormous impact on a large range of fields such as computer vision, robotics and computational biology. These methods have allowed researchers to exploit evidence from data to learn models in a principled manner. One of the most important developments has been that of kernel methods

This work was supported by the EU project TOMSY (IST-FP7-270436) as well as the Swedish Foundation for Strategic Research.

An earlier version of this work was published under the title “The Path Kernel” at the International Conference on Pattern Recognition Applications and Methods, 2013.

A. Baisero (✉) · F.T. Pokorny · D. Kragic · C.H. Ek
Computer Vision and Active Perception Laboratory, Centre for Autonomous Systems,
KTH Royal Institute of Technology, Stockholm, Sweden
e-mail: baisero@kth.se

F.T. Pokorny
e-mail: fpokorny@kth.se

D. Kragic
e-mail: danik@kth.se

C.H. Ek
e-mail: chek@kth.se

[5] which embed the input data in a potentially high-dimensional vector space with the intention of achieving improved robustness of classification and regression techniques. The main benefit of kernel methods is that, rather than defining an explicit feature space that has the desired properties, the embedding is characterised implicitly through the choice of a kernel function which models the inner product in an induced space. This creates a very natural paradigm for recovering the desired characteristics of a representation. Kernel functions based on a stationary distances (usually an \mathcal{L}_p – norm) have been particularly successful in this context [3]. However, for many application domains, the data does not naturally lend itself to a finite dimensional vectorial representation. Symbolic sequences and graphs, for example, pose a problem for such kernels.

For non-vectorial data, the techniques used for learning and inference are generally much less developed. A desirable approach is hence to first place the data in a vector space where the whole range of powerful machine learning algorithms can be applied. Simple approaches such as the Bag-of-Words model, which creates a vectorial representation based on occurrence counts of specific representative “words”, have had a big impact on computer vision [18]. These methods incorporate the fact that a distance in the observed space of image features does not necessarily reflect a similarity between the observed scenes. Another approach, where strings are transformed into a vectorial representation before a kernel method is applied, has been the development of string kernels [14, 17]. Such kernels open up a whole range of powerful techniques for non-vectorial data and they have been applied successfully to robotics [15], computer vision [13] and biology [12]. Other related works are based on convolution kernels [11]. Using such kernels, a vectorial representation that respects the structure of a graph can be recovered. Another approach to define an inner product between sequences is to search for a space where similarity is reflected by “how well” sequences align [6, 7, 19].

In this paper, we present a new kernel for representing sequences of symbols which extends and further develops the concept of sequence alignment. Our kernel is based on a *ground space* which encodes similarities between the symbols in a sequence. We show that our kernel is a significant improvement compared to the state of the art both in terms of computational complexity and in terms of its ability to represent the data.

2 Kernels and Sequences

Before we proceed with describing previous work for creating kernel induced feature spaces for sequences, we will clarify our notation and our notion of kernels. When discussing kernels in the context of machine learning, we have to distinguish between several uses of the word kernel. In this paper, a kernel denotes any symmetric function $k : X \times X \rightarrow \mathbb{R}$, where X is a non-empty set [10]. A positive semi-definite (psd) kernel is a kernel $k : X \times X \rightarrow \mathbb{R}$ such that $\sum_{i,j}^n c_i c_j k(x_i, x_j) \geq 0$ for any $\{x_1, \dots, x_n\} \subset X$, $n \in \mathbb{N}$ and $c_1, \dots, c_n \in \mathbb{R}$. If the previous inequality is strict

when $c_i \neq 0$ for at least one $i \in \{1, \dots, n\}$, the kernel is called positive definite (pd). Further specialisations, such as negative definite (nd) kernels, exist and are of independent interest.

While there are strong theoretical results on the existence of embeddings corresponding to psd kernels [2, p. 22], non-psd kernel functions can still be useful in applications. Examples of kernels that are known to be neither pd nor psd but which are still successfully used in classification include [9]. On another note, there are also kernels which are conjectured to be psd, and which have been shown to be psd in experiments, but for which there currently is no proof for the corresponding positiveness [1].

In this work, we consider finite sequences of symbols belonging to an alphabet set Σ , i.e. $s = (s_1, s_2, \dots, s_{|s|})$ denotes such a sequence, with $s_i \in \Sigma$, and where $|s| \in \mathbb{N}_0$ denotes the length of the sequence. We denote by $s_{a:b}$, with $1 \leq a < b \leq |s|$, the subsequence $s_{a:b} = (s_a, \dots, s_b)$. When the indices a or b are omitted, they implicitly refer to 1 or $|s|$ respectively. The *inverse* of a sequence s is defined by $\text{inv}(s)_i = s_{|s|-(i-1)}$.

In this work, we assume that we are given a psd kernel function $k_\Sigma : \Sigma \times \Sigma \rightarrow \mathbb{R}$ describing the similarity between elements of the alphabet Σ and will refer to k_Σ as the ground kernel. Given k_Σ , we can now define the *path matrix*.

Definition 1 (Path Matrix) Given two finite sequences s, t with elements in an alphabet set Σ and a kernel $k_\Sigma : \Sigma \times \Sigma \rightarrow \mathbb{R}$, we define the *path matrix* $G(s, t) \in \mathbb{R}^{|s| \times |t|}$ by $[G(s, t)]_{ij} = k_\Sigma(s_i, t_j)$.

We denote $\delta_{00} \stackrel{\text{def}}{=} (0, 0)$, $\delta_{10} \stackrel{\text{def}}{=} (1, 0)$, $\delta_{01} \stackrel{\text{def}}{=} (0, 1)$, $\delta_{11} \stackrel{\text{def}}{=} (1, 1)$ and $S \stackrel{\text{def}}{=} \{\delta_{10}, \delta_{01}, \delta_{11}\}$. S is called the set of admissible steps. A sequence of admissible steps starting from $(1, 1)$ defines the notion of a path:

Definition 2 (Path) A path over a $m \times n$ path-matrix G is a map $\gamma : \{1, \dots, |\gamma|\} \rightarrow \mathbb{N} \times \mathbb{N}$ such that

$$\gamma(1) = (1, 1), \quad (1)$$

$$\gamma(i+1) = \gamma(i) + \delta_i, \quad \text{for } 1 \leq i < |\gamma|, \text{ with } \delta_i \in S, \quad (2)$$

$$\gamma(|\gamma|) = (m, n). \quad (3)$$

$|\gamma|$ and δ_i denote the path's *length* and *i*th *step* respectively. Furthermore, we adopt the notation $\gamma(i) = (\gamma_X(i), \gamma_Y(i))$. A path determines *stretches*, or alignments, on the input sequences according to $s_{\gamma_X} = (s_{\gamma_X(1)}, \dots, s_{\gamma_X(|\gamma|)})$ and $t_{\gamma_Y} = (t_{\gamma_Y(1)}, \dots, t_{\gamma_Y(|\gamma|)})$.

We denote the set of all paths on a $m \times n$ matrix as $\Gamma(m, n)$. Its cardinality is equal to the Delannoy number $Del(m, n)$.

2.1 Sequence Similarity Measures

A popular similarity measure between time-series is Dynamic Time Warping (DTW) [8, 16], which determines the distance between two sequences s and t as the minimal score obtained by all paths, i.e.

$$d_{\text{DTW}}(s, t) = \min_{\gamma \in \Gamma} D_{s,t}(\gamma), \quad (4)$$

where $D_{s,t}$ represents the score of a path γ defined by

$$D_{s,t}(\gamma) = \sum_{i=1}^{|\gamma|} \varphi(s_{\gamma_X(i)}, t_{\gamma_Y(i)}), \quad (5)$$

where φ is some given similarity measure. However, DTW lacks a geometrical interpretation in the sense that it does not necessarily respect the triangle inequality [7]. Furthermore, this similarity measure is not likely to be robust as it only uses information from the minimal cost alignment.

Taking the above into consideration, Cuturi et al. suggest a kernel referred to as the *Global Alignment Kernel* [7]. Instead of considering the minimum over all paths, the Global Alignment Kernel combines all possible path scores. The kernel makes use of an exponentiated soft-minimum of all scores, generating a more robust result which reflects the contents of all possible paths:

$$k_{\text{GA}}(s, t) = \sum_{\gamma \in \Gamma} e^{-D_{s,t}(\gamma)}. \quad (6)$$

By taking the ground kernel to be $k_{\Sigma}(\alpha, \beta) = e^{-\varphi(\alpha, \beta)}$, k_{GA} can be described using the path matrix as

$$k_{\text{GA}}(s, t) = \sum_{\gamma \in \Gamma} \prod_{i=1}^{|\gamma|} G(s, t)_{\gamma(i)}. \quad (7)$$

The leading principle in this approach is hence a combination of kernels on the level of symbols over all paths along $G(s, t)$. Cuturi shows that incorporating all the elements of $G(s, t)$ into the final results can vastly improve classification compared to using only the minimal cost path. Furthermore, k_{GA} is proven to be psd under the condition that both k_{Σ} and $\frac{k_{\Sigma}}{1+k_{\Sigma}}$ are psd [7], giving foundation to its geometrical interpretation.

However, Cuturi's kernel makes use of products between ground kernel evaluations along a path. This implies that the score for a complete path will be very small if $\varphi(s_i, t_j)$ is sufficiently large, which leads to the problem of diagonally dominant kernel matrices [6, 8] from which the global alignment kernel suffers. The issue is particularly troubling when occurring at positions near the top-left or bottom-right

corners of the path matrix because it will affect many of the paths. Furthermore, paths contribute with equal weight to the value of the kernel. To reduce this effect, it is suggested in [7] to rescale the kernel values and use its logarithm instead. We argue that paths which travel closest to the main diagonal of the path matrix should be considered as more important than others, since they minimise the distortion imposed on the input sequences, i.e. s_{γ_X} and t_{γ_Y} are then most similar to s and t . To rectify this and to include a preference towards diagonal paths, a generalisation called the *Triangular Global Alignment Kernel* was developed, which considers only a subset of the paths [6]. This generalisation imposes a crude preference for paths which do not drift far away from the main diagonal.

In this paper, we develop a different approach by introducing a weighting of the paths in Γ based on the number of diagonal and off-diagonal steps taken. We manipulate the weights to encode a preference towards consecutive diagonal steps while at the same time accumulating information about all possible paths. Furthermore, by replacing the accumulation of symbol kernel responses along the path using a summation rather than a product, the kernel’s evaluation reflects more gracefully the structure of the sequences and avoids abrupt changes in value.

3 The Path Kernel

In this section we will describe our proposed kernel which we will refer to as the *path kernel*. Figure 1 illustrates the contents of a path matrix in a simplified example, showing the emergence of diagonal patterns when the two sequences are in good correspondence.

Table 1 shows the resulting alignments associated with the paths shown in Fig. 1. We argue that the values, the length and the location of these diagonals positively reflect the relation between the inputs and should thus be considered in the formu-

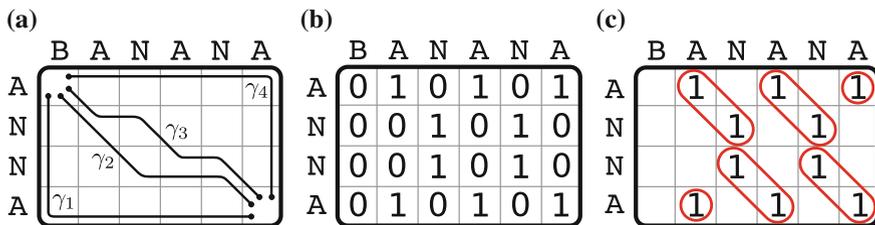


Fig. 1 Illustration of the concept of paths and the contents of $G(s, t)$ for $s = \text{“ANNA”}$ and $t = \text{“BANANA”}$. On the *left*, we illustrate a small number of paths which traverse G . The path kernel makes use of these, together with all the other paths, to collect data from the matrix and to extract a final score. In the *center*, we display the contents of $G(s, t)$, assuming $k_{\Sigma}(\alpha, \beta) = \delta_{\alpha\beta}$, i.e. Kronecker’s delta function. On the *right*, we highlight the corresponding diagonals whose number, length and position are related to the similarity between subsequences of s and t

Table 1 Stretches associated to the paths in Fig. 1a with the underlined substrings denoting a repeated symbol

| Stretches | | | |
|------------|--------------------------|------------|--------------------------|
| γ_1 | <u>A N N A A A A A A</u> | γ_2 | <u>A N N N N A</u> |
| | <u>B B B A N A N A</u> | | <u>B A N A N A</u> |
| γ_3 | <u>A N N N N A</u> | γ_4 | <u>A A A A A N N A</u> |
| | <u>B A N A N A</u> | | <u>B A N A N A A A A</u> |

Note that, even though γ_2 and γ_3 produce the same stretches, they traverse the matrix differently and should thus be considered separately

lation of a good kernel. High values imply a good match on the ground kernel level, while their length encodes the extent of the match. On the other hand, the position relative to the main diagonal reflects how much the input sequences had to be “stretched” in order for the match to be encountered. We wish to have a feature space where a smaller stretch implies a better correspondence between the sequences.

Let us now define a new kernel that incorporates different weightings depending on the steps used to travel along a path.

Definition 3 (*Path Kernel*) For any sequences s, t , we define

$$k_{\text{PATH}}(s, t) \stackrel{\text{def}}{=} \begin{cases} k_{\Sigma}(s_1, t_1) \\ \quad + C_{\text{HV}} k_{\text{PATH}}(s_2:, t) \\ \quad + C_{\text{HV}} k_{\text{PATH}}(s, t_2:) \\ \quad + C_{\text{D}} k_{\text{PATH}}(s_2:, t_2:) \\ 0 \end{cases} \quad \begin{array}{l} |s| \geq 1 \wedge |t| \geq 1, \\ \\ \\ \text{otherwise,} \end{array} \quad (8)$$

where C_{HV} and C_{D} represent weights assigned to vertical or horizontal steps and diagonal steps respectively. By enforcing the constraints $C_{\text{HV}} > 0$ and $C_{\text{D}} > C_{\text{HV}}$, we aim to increase the relative importance of paths with many diagonal steps.

The symmetry of the kernel is easily verifiable. On the other hand, the positive semi-definiteness of the kernel is not immediately obvious from the definition.¹

3.1 Efficient Computation

Kernel methods often require the computation of a kernel function on a large dataset, where the number of kernel evaluations will grow quadratically with the number of data-points. It is hence essential that the kernel evaluations themselves are efficiently computable.

Not only can the path kernel be evaluated using a Dynamic Programming algorithm which avoids the expensive recursion in (8) and which achieves a computa-

¹ In an extension of this work, which is currently under review, we provide a proof of the positive semi-definiteness of our kernel.

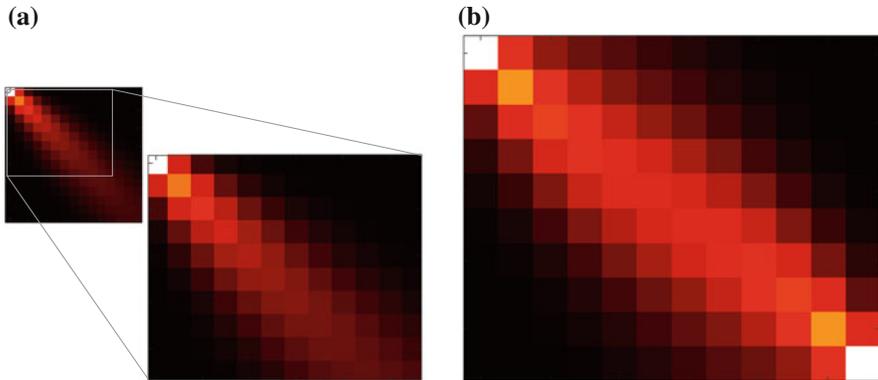


Fig. 2 On the *left*, a precomputed 15×15 weight matrix with $C_{HV} = 0.3$ and $C_D = 0.34$ is used to select a 10×12 weight matrix which can then be used to evaluate $k_{\text{PATH}}(s, t)$ for input sizes $|s| = 10$ and $|t| = 12$. On the *right*, the inversion invariant $\tilde{\omega}_{\text{PATH}}$ corresponding to \tilde{k}_{PATH} for the same input sizes is displayed

tional complexity comparable with DTW and k_{GA} , but it can also be computed very efficiently using the following alternative formulation:

$$k_{\text{PATH}}(s, t) = \sum_{ij} G(s, t)_{ij} \omega_{\text{PATH}ij}, \tag{9}$$

$$[\omega_{\text{PATH}}]_{ij} = \sum_{d=0}^{\min(i,j)-1} C_{\text{HV}}^{i+j-2-2d} C_D^d (d, i-1-d, j-1-d)!. \tag{10}$$

The usefulness of (9) comes from the fact that the contents of the weight matrix ω_{PATH} do not really depend on s, t , and thus ω_{PATH} can in fact be pre-computed up to an adequate size² (Fig. 2). After this, the evaluation of the kernel for input of sizes m and n is achieved by simply selecting the sub-matrix ranging from $(1, 1)$ to (m, n) ; the remaining matrix element-wise multiplication can then be efficiently carried out. By taking advantage of this, one can evaluate the kernel at speeds depending only on the speed of the evaluation of G and the speed of a simple matrix multiplication (with the initial overhead consisting of the pre-computation of ω_{PATH}). The weight matrix can also be computed through an efficient and very simple Dynamic Programming algorithm similar to the one which can be used to evaluate the kernel itself.

We call a kernel satisfying $k(s, t) = k(\text{inv}(s), \text{inv}(t))$ inversion invariant. If a kernel k does not naturally have this property, it can be enforced by replacing k with

$$\tilde{k}(s, t) = \frac{1}{2} [k(s, t) + k(\text{inv}(s), \text{inv}(t))]. \tag{11}$$

² For any specific dataset, that would be the length of the longest sequence.

The path kernel is not originally inversion invariant, but invariance can be enforced without the need for a double computation of the kernel for each evaluation. This is done by modifying the selected sub-matrix of ω_{PATH} as follows: for any two inputs with lengths m and n , we replace the weight matrix ω_{PATH} by

$$[\tilde{\omega}_{\text{PATH}}]_{ij} = \frac{1}{2} [\omega_{\text{PATH}ij} + \omega_{\text{PATH}m-i+1, n-j+1}], \quad (12)$$

and then proceed using this weight matrix.

3.2 Ground Kernel Choice

The path kernel is based on a ground kernel which, apart from being a psd kernel function, is not constrained in any other way. However, we show in this paragraph that an arbitrarily k_{Σ} may lead to undesirable results.

Assume an alphabet and a ground kernel such that $\alpha, \beta \in \Sigma$, $k_{\Sigma}(\alpha, \alpha) = 1$, $k_{\Sigma}(\beta, \beta) = 1$ and $k_{\Sigma}(\alpha, \beta) = -1$. Given the input sequences $s = (\alpha, \beta, \dots, \alpha, \beta)$ and $t = (\beta, \alpha, \dots, \beta, \alpha)$, one may be inclined to say that s and t are very similar because each can be obtained from the other by cyclically shifting the symbols by one position. However, the contents of $G(s, t)$ show a collection of ones and negative ones organised in a chessboard-like disposition. This obviously leads to heavy fluctuations during the computation of $k_{\text{PATH}}(s, t)$ and to potentially very small values. Furthermore, the issue is present even in the computation of $k_{\text{PATH}}(s, s)$ and $k_{\text{PATH}}(t, t)$ which is not desirable under any circumstance. This problem is however easily rectifiable by considering only ground kernels that yield non-negative results on elements of Σ .

4 Experiments

In this section, we present the results of experiments performed with the proposed kernel. In particular, we perform two separate quantitative experiments that, in addition to our qualitative results, shed some light on the behaviour of the proposed method in comparison to previous work. We will compare our approach to k_{GA} as well as the non-psd kernel obtained by using the negative exponential of the DTW distance,

$$k_{\text{DTW}}(s, t) = e^{-d_{\text{DTW}}(s, t)}. \quad (13)$$

In the first experiment, we generate eight different classes of uni-variate sequences. Each class consists of a periodic waveform namely $\pm\text{sine}$, $\pm\text{cosine}$, $\pm\text{sawtooth}$ and $\pm\text{square}$. From these, we generate noisy versions by performing three different operations:

1. The length of the sequence is generated by sampling from a normal distribution $\mathcal{N}(100, \sigma_l^2)$, rounding the result and rejecting non-positive lengths.
2. We obtain an *input* sequence as $|s|$ equidistant numbers spanning 2 periods of the wave; we then add to each element an *input* noise which follows a normal distribution $\mathcal{N}(0, \sigma_l^2)$.
3. We feed the noisy input sequence to the generating waveform and get an *output* sequence, to which we add *output* noise which follows a normal distribution $\mathcal{N}(0, \sigma_o^2)$.

Figure 3 shows the sequences for the parameter setting $\sigma_{\{l,i,o\}} = 5$. This corresponds to the setting which we will use to present our main results.

The path kernel has two different sets of parameters: the ground kernel and the weights associated with steps in the path matrix. In our experiments, we use a simple zero mean Gaussian kernel with standard deviation 0.1 as ground kernel. The step weights C_{HV} and C_D are set to 0.3 and 0.34 respectively. We use the same setting throughout the experiments. The behaviour of the kernel will change with the value of these parameters. A complete analysis of this is however beyond the scope of this paper. Here, we focus on the general characteristics of our kernel which summarises all possible paths using step weights satisfying $C_{HV} < C_D$ —implying a preference for diagonal paths.

In order to get an understanding of the geometric configuration of the data that our kernel matrices corresponds to, we project the data onto the two first principal directions as determined by each of these kernels. The result can be seen in Fig. 4. It is important to note that, as the DTW kernel has negative eigenvalues, it does not imply a geometrically valid configuration of datapoints in a feature space.

From Fig. 4, we get a qualitative understanding of how the induced feature spaces looks like. However, a representation is simply the means to an end and to be able to make a valuable assessment of its useability, we need to use it to achieve a task. We do so through two different experimental setups: The first is meant to test the discriminative capabilities of the representation; the second evaluates how well the representation is suited for generalisation.

In order to test the discriminability of the feature space generated by the path kernel, we perform a classification experiment using the same data as explained above, and where the task is to predict the generating class of a waveform. We feed the kernel matrix into an SVM classifier [4], use a 2-fold cross-validation, and report the average over 50 runs. Due to the negative eigenvalues, the classification fails for the DTW kernel. For this reason, we only present results for the remaining two kernels. In Fig. 5a, the results for the classification with increasing noise levels are shown. For moderate noise-levels (up to $\sigma_{\{l,i,o\}} = 5$), the global alignment and the path kernel are comparable in performance, while—at a higher noise level—the performance of the global alignment kernel rapidly deteriorates and at $\sigma_{\{l,i,o\}} = 9$ its performance is about chance, while the path kernel still achieves a classification rate of over 80 %.

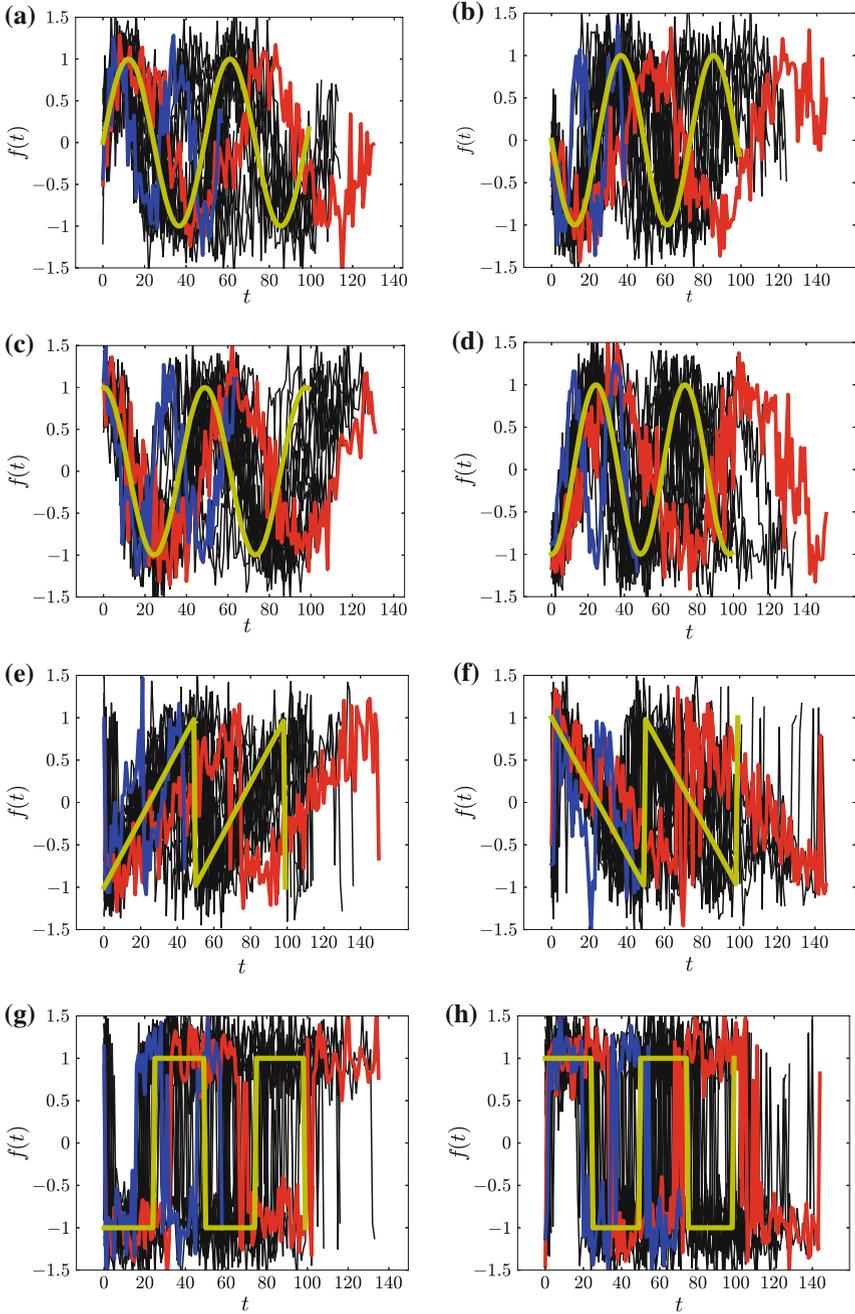


Fig. 3 The figure above shows the eight different waveforms used for the classification for a noise level corresponding to $\sigma_{\{l,i,o\}} = 5$. The golden curve depicts the base waveform without noise while the blue and red curves show the shortest and the longest noisy example respectively. The black curves display the remaining examples in the database

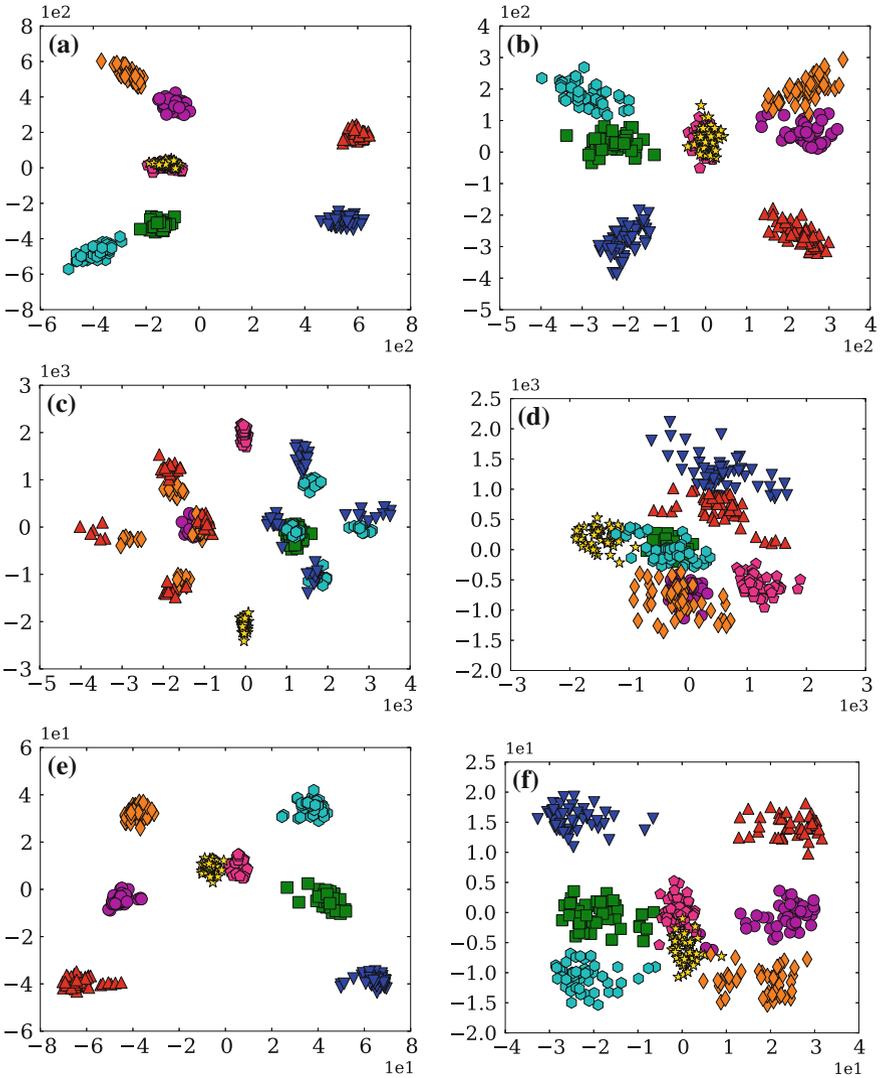


Fig. 4 The above figure displays the two dimensional principal subspace for the DTW Kernel (*top*), Global Alignment Kernel (*middle*) and the Path Kernel (*bottom*). The left column represents data with generation noise $\sigma_{\{l,i,o\}} = 2$, while, in the *right column*, the noise is increased to $\sigma_{\{l,i,o\}} = 5$, corresponding to the waveforms in Fig. 3. The different waveforms are displayed as follows: sine and-sine as a *magenta circle* and a *green square*, cosine and-cosine as a *pink pentagon* and a *yellow star*, sawtooth and-sawtooth as a *light-blue hexagon* and an *orange diamond* and *square and-square* as a *blue* and a *red triangle* respectively

The classification experiment shows that the path kernel significantly outperformed the global alignment kernel when noise in the sequence became significant. Looking at the feature space, depicted in Fig. 4, we see that the path kernel encodes a feature space having more clearly defined clusters corresponding to the different waveforms. Additionally, the clusters also have a simpler structure. This signifies that the path kernel should be better suited for generalisation purposes, where it is beneficial to have a large continuous region of support which gracefully describes the variations in the data—rather than working in a space that barely separates the classes.

We now generate a new dataset consisting of 100 noisy sine-waves ($\sigma_{l,i,o} = 5$) shifted in phase between 0 and π . The data is split uniformly into two halves and the first is used for training and the second for testing. We want to evaluate how well the kernel is capable of generalising over the training data. To that end we regress from the proportion of the training data to the test data and evaluate how the prediction error changes by altering this proportion. The prediction is performed using simple least-square regression over the kernel induced feature space. Figure 5b shows the results using different sizes of the training data; The path kernel performs significantly better compared to the global alignment kernel and the results improve with the size of the training dataset. Interestingly, the global alignment kernel produces very different results dependent on the size of the training dataset indicating that it is severely over-fitting the data.

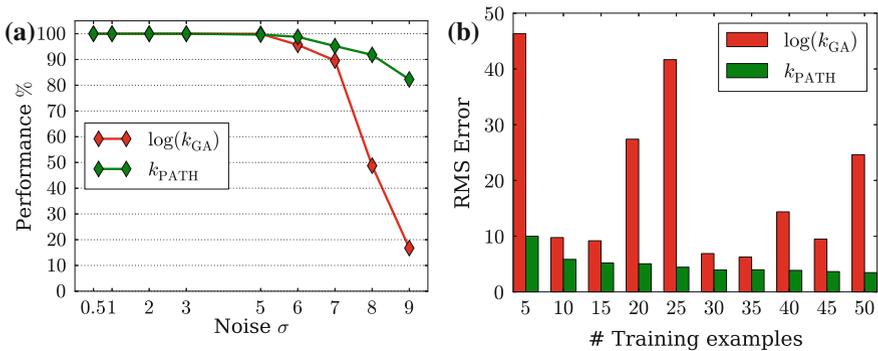


Fig. 5 The *left* figure displays the classification rate for predicting the waveform type using an SVM classifier in the feature space defined by the Global Alignment Kernel (*blue*) and the Path Kernel (*green*). The *x-axis* depicts the noise level parametrized by $\sigma_{l,i,o}$. The *right* figure depicts the RMS error when predicting the phase shift from a noisy sine waveform by a regression over the feature space induced by the kernels. The *red bars* correspond to the Global Alignment Kernel and the *green bars* to the Path Kernel. The *y-axis* shows the error in percentage of phase, while the *x-axis* indicates the size of the training dataset. The test set has a constant size of 50

5 Conclusions

In this paper, we have presented a novel kernel for encoding sequences. Our kernel reflects and encodes all possible alignments between two sequences by associating a cost to each. This cost encodes a preference towards specific paths. The kernel is applicable to any kind of symbolic or numerical data as it requires only the existence of a kernel between symbols. We have presented both qualitative and quantitative experiments exemplifying the benefits of the path kernel compared to competing methods. We show that the proposed method significantly improves results both with respect to discrimination and generalisation especially in noisy scenarios. The computational cost associated with the kernel is considerably lower than competing methods, making it applicable to data-sets that could previously not be investigated using kernels.

In this paper, we have chosen a very simple dataset in order to evaluate our kernel. Given our encouraging results, we are currently working on applying our kernel to more challenging real-world datasets. Additionally, we are investigating the possibility of optimizing the kernel parameters to further improve classification results.

References

1. Bahlmann, C., Haasdonk, B., Burkhardt, H.: Online handwriting recognition with support vector machines—a kernel approach. In: 8th International Workshop on Frontiers in Handwriting Recognition, pp. 49–54 (2002)
2. Berlinet, A., Thomas-Agnan, C.: Reproducing Kernel Hilbert Spaces in Probability and Statistics. Kluwer, Dordrecht (2004)
3. Buhmann, M.D., Martin, D.: Radial Basis Functions: Theory and Implementations. Cambridge University Press, Cambridge (2003)
4. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines (2001)
5. Cristianini, N., Shawe-Taylor, J.: An Introduction to Support Vector Machines: And Other Kernel-Based Learning Methods. Cambridge University Press, Cambridge (2006)
6. Cuturi, M.: Fast global alignment kernels. In: Proceedings of the International Conference on Machine Learning (2010)
7. Cuturi, M., Vert, J.P., Birkenes, O., Matsui, T.: A kernel for time series based on global alignments. In: IEEE International Conference on Acoustics, Speech and Signal Processing, pp. II-413–II-416 (2007)
8. Gudmundsson, S., Runarsson, T.P., Sigurdsson, S.: Support vector machines and dynamic time warping for time series. In: Proceedings of the IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence) (x), pp. 2772–2776 (2008). <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4634188>
9. Haasdonk, B.: Feature space interpretation of SVMs with indefinite kernels. IEEE Trans. Pattern Anal. Mach. Intell. **27**(4), 482–492 (2005)
10. Haasdonk, B., Burkhardt, H.: Invariant kernel functions for pattern analysis and machine learning. Mach. learn. **68**(1), 35–61 (2007)
11. Haussler, D.: Convolution kernels on discrete structures. Tech. rep. (1999)
12. Leslie, C., Kuang, R.: Fast string kernels using inexact matching for protein sequences. J. Mach. Learn. Res. **5**, 1435–1455 (2004)

13. Li, M., Zhu, Y.: Image classification via LZ78 based string kernel: a comparative study. *Advances in Knowledge Discovery and Data Mining*. Springer, Berlin (2006)
14. Lodhi, H., Saunders, C., Shawe-Taylor, J., Cristianini, N., Watkins, C.: Text classification using string kernels. *J. Mach. Learn. Res.* **2**, 419–444 (2002)
15. Luo, G., Bergström, N., Ek, C.H., Kragic, D.: Representing actions with Kernels. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2028–2035 (2011)
16. Sakoe, H., Chiba, S.: Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Acoust. Speech Signal Process.* **26**(1), 43–49 (1978)
17. Saunders, C., Tschach, H., Shawe-Taylor, J.: Syllables and other string kernel extensions. In: *Proceedings of the Nineteenth International Conference on Machine Learning (ICML'02)* (2002)
18. Sivic, J., Zisserman, A.: Efficient visual search of videos cast as text retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(4), 591–606 (2009)
19. Watkins, C.: Dynamic alignment kernels. In: *Proceedings of the Advances in Neural Information Processing Systems* (1999)