

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/175286>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

MAXIMUM MEAN DISCREPANCY KERNELS FOR PREDICTIVE AND PROGNOSTIC MODELING OF WHOLE SLIDE IMAGES

*Piotr Keller**, *Muhammad Dawood** and *Fayyaz ul Amir Afsar Minhas*

Tissue Image Analytics Center, University of Warwick, United Kingdom

ABSTRACT

How similar are two images? In computational pathology, where Whole Slide Images (WSIs) of digitally scanned tissue samples from patients can be multi-gigapixels in size, determination of degree of similarity between two WSIs is a challenging task with a number of practical applications. In this work, we explore a novel strategy based on kernelized Maximum Mean Discrepancy (MMD) analysis for determination of pairwise similarity between WSIs. The proposed approach works by calculating MMD between two WSIs using kernels over deep features of image patches. This allows representation of an entire dataset of WSIs as a kernel matrix for WSI level clustering, weakly-supervised prediction of TP53 mutation status in breast cancer patients from their routine WSIs as well as survival analysis with state of the art prediction performance. We believe that this work will open up further avenues for application of WSI-level kernels for predictive and prognostic tasks in computational pathology.

Index Terms— Computational Pathology, Weakly Supervised Learning, Survival Analysis, Mutation Prediction

1. INTRODUCTION

Computational Pathology (CPath) involves predictive modeling of routine histology Whole Slide Images (WSIs) for various diagnostic and prognostic tasks such as detection of cancer, discovery of image patterns associated with various genetic mutations and survival analysis which are conventionally, performed using deep learning approaches [1, 2]. Due to the sheer size of multi-gigapixel WSIs and the associated memory requirements of such predictive problems, WSIs are divided into patches and passed through a neural network to generate patch level prediction scores. These scores are then aggregated to generate WSI-level predictions. However, this strategy is inherently unable to capture information distributed over scale greater than the patch size. To capture higher level structural information, we need to analyze larger regions or perform different types of aggregation ranging from simple averaging or max-pooling [3] to more sophisticated hierarchical attention models [4] and graph neural networks [5].

In this work, we approach the aggregation problem from a different perspective - the use of Maximum Mean Discrepancy (MMD) as a measure of (dis)similarity between WSIs (see Fig. 1). MMD is rooted in kernel matrix theory and is used as a statistical test to determine if two sets of measurements are drawn from different distributions or not [6]. The primary hypothesis underlying this work is that (dis)similarity between two WSIs can be captured by their MMD by simply representing each WSI as a set of features of its constituent patches. This not only enables computation of pairwise WSI (dis)similarity, the resulting metric also leads to formulation of WSI-level kernels thus opening up the possibility for the use of the existing repertoire of powerful and versatile kernel methods for predictive modeling. Such kernels can capture slide level similarity in a dataset of WSIs and thus can be used for clustering, classification or survival prediction as well as integration of multi-modal data. To the best of our knowledge, this is the first report of WSI level kernels in this domain and it paves the way for using such methodologies for other related tasks.

Our primary technical contribution in this work is the proposed use of WSI-level MMD kernels for solving both unsupervised classification problems such as clustering and data mining as well as weakly-supervised predictive and prognostic tasks in CPath in which training label or target information is only available at the WSI-level. As a proof of principle, we use this approach to solve two clinically relevant problems - prediction of TP53 mutation status and survival analysis of breast cancer patients using their routine Hematoxylin and Eosin (H&E) stained tissue slides [7]. Below, we introduce both of these problems along with a brief description of existing work in this domain.

In the context of breast cancer diagnosis, detection of mutations in tumour suppressor genes such as TP53 is of prime importance [8]. TP53 mutations, found in 20–40% of all breast cancer cases, can be determinants of underlying cancer pathways and treatment options [9, 10]. Previous work has shown that while it is possible to predict TP53 mutation status from WSIs, the accuracy of such predictions is quite low with state of the art area under the receiver operating characteristic curve (AUC-ROC) of around 0.72 [1].

Another important yet challenging problem in CPath is survival analysis which aims to analyze the expected duration

* Joint first authorship

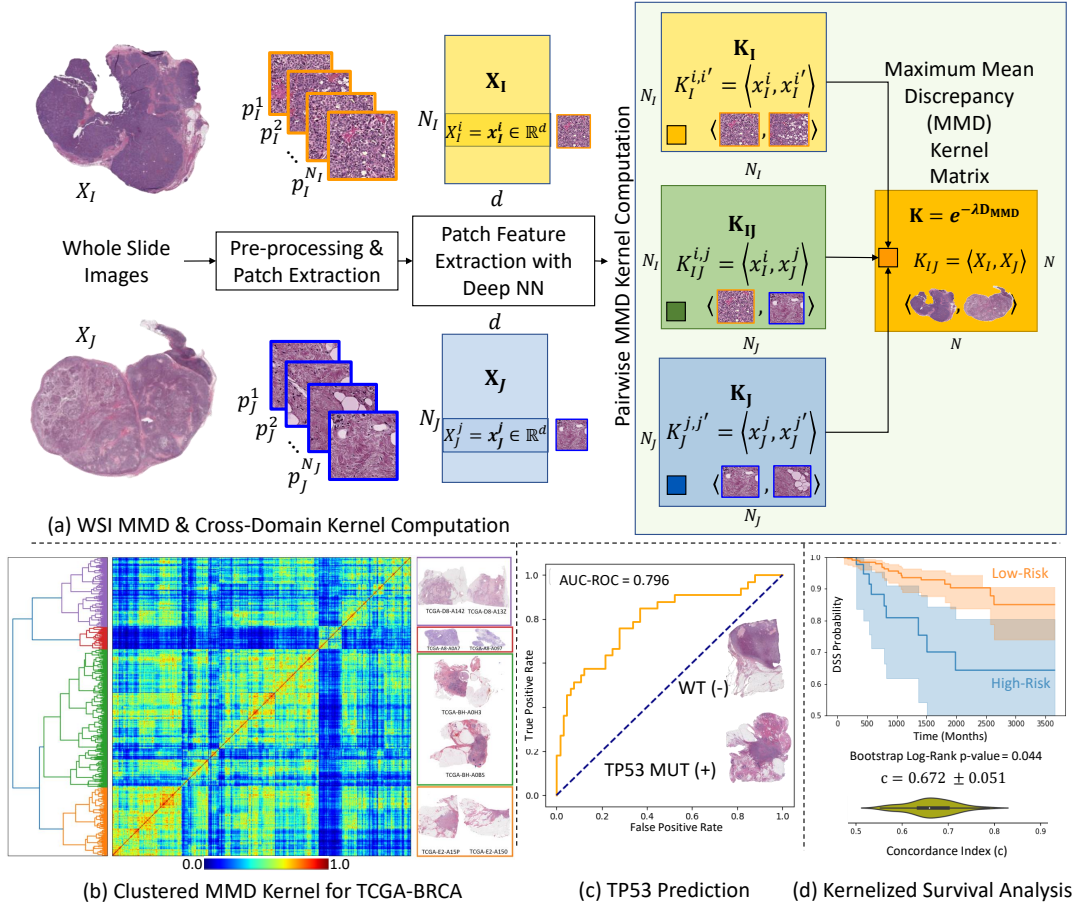


Fig. 1: Proposed workflow and key results. (a) WSIs are tiled into non-overlapping image patches. We then encode each patch into 1024-dimensional feature vector by extracting latent representation from a pre-trained ShuffleNet. The novel WSI-level Maximum Mean Discrepancy matrix and the corresponding MMD Kernel is computed (Eq. 3). (b) The resulting MMD kernel for TCGA-BRCA dataset with hierarchical clustering to illustrate WSI-level groupings within the dataset. Results of TP53 mutation status prediction (in (c)) and survival analysis (in (d)) are shown using MMD kernel of TCGA-BRCA dataset.

of time until an event such as progression or death occurs. [11] Current state of the art models for survival prediction from routine WSI data achieve a relatively low Concordance Index (C-Index) of no more than 0.58 [12, 2, 13].

2. MATERIALS AND METHOD

2.1. Study dataset

We collected 1,133 publicly available WSIs of Formalin-Fixed paraffin-Embedded (FFPE) Hematoxylin and Eosin (H&E) stained tissue section of 1084 breast cancer patients from The Cancer Genome Atlas [14, 15] (TCGA-BRCA). In line with the previous study with state of the art results for point mutation prediction [1], we filtered out WSIs 1) with poor visual quality; (2) abnormal staining 3) with low informative tissue region; and (4) missing baseline resolution information. After slide selection, we end-up with 652 WSIs belonging to 652 patients. The TP53 mutation status

data for these patients was downloaded from CBioPortal (417 negative, 235 positive) [16, 17].

For survival analysis, we used the same dataset. However, previous state of the art approaches in this domain had used the whole dataset for analysis. In order to have a meaningful comparison, we have used the same cohort of patients. The Disease Specific Survival (DSS) data for these patients were collected from the TCGA Pan-Cancer Clinical Data Resource (TCGA-CDR) [18]. Patients with missing DSS data was not used in the study. In total we used WSI of 1,033 patients in our study.

2.2. Pre-processing and Representation of WSIs

Let $D = \{X_1, X_2, \dots, X_N\}$ be the set of N WSIs in the dataset. We segment the tissue region of a WSIs using a tissue segmentation model to remove artefacts such as pen markings, tissue folding, etc. As an entire WSI at full resolution can be very large in size ($100,000 \times 150,000$ pixels) and does

not fit into a GPU memory, we tile each WSI into patches of size 512×512 at a spatial resolution of 0.50 microns-per-pixel (MPP). Patches capturing less than 40% of informative tissue area (mean pixel intensity above 200) are discarded, and the rest of patches (both tumour, and non-tumour) are used. We model a given WSI (indexed by I) as a set of patches as $X_I = \{x_I^1, x_I^2, \dots, x_I^{N_I}\}$ such that N_I is the number of patches in the image and x_I^i represents the i th patch in image X_I . For each patch, we obtain its $d = 1024$ -dimensional feature representation $x_i^I \in R^d$ by passing it through a convolutional neural network (CNN) encoder. Any feature extractor can be used as patch level encoder. However for this work, we extracted the latent representation of the second-last fully connected layer of an ImageNet pretrained ShuffleNet [19].

2.3. Maximum Mean Discrepancy Kernels between WSIs

Modelling WSIs as sets enables use of set-based discrepancy metrics such as MMD for measuring differences between any two WSIs. Originally proposed as a kernel two-sample test [6], MMD between two WSIs can be defined as [20]:

$$MMD(X_I, X_J) = \|\mathbb{E}_{x_i^i \sim X_I}[\varphi(x_i^i)] - \mathbb{E}_{x_j^j \sim X_J}[\varphi(x_j^j)]\|_{\mathcal{H}} \quad (1)$$

Here, we assume that the WSIs X_I and X_J come from some space \mathcal{X} . We also have an abstract function $\varphi: \mathcal{X} \mapsto \mathcal{H}$ and \mathcal{H} is the RKHS that allows us to perform the change of basis on the patches. Thus, MMD measures discrepancy in various moments of feature distributions of the patches in the two WSIs. Specifically, defining $\varphi(x) = x$ will result in the MMD between the two WSIs being simply the discrepancy between average patch representations of the two WSIs: $MMD(X_I, X_J) = \|\mathbb{E}_{x_i^i \sim X_I}[x_i^i] - \mathbb{E}_{x_j^j \sim X_J}[x_j^j]\|_{R^d} = \|\mu_I - \mu_J\|$. Similarly, defining $\varphi(x) = (x, x^2)$ enables the resulting MMD to capture differences in means (as with $\varphi(x) = x$) as well as variances (second order moments) of corresponding features. It is important to emphasize that even though the approach is called Maximum "Mean" Discrepancy, through the use of kernels, it can calculate arbitrary higher order differences as discussed below.

By varying our definition of φ we can compare an arbitrary WSI-level statistical moment. However, it is often the case that two different distributions share some moments and thus may be seen as identical by MMD when in reality they are not. To combat this we would ideally like to compare an infinite number of moments which can be done by using kernels as these allow for implicit calculations of moments. To achieve this, we use the patch-level Gaussian kernel defined as: $k(x_I^i, x_J^j) = e^{-\frac{\|x_I^i - x_J^j\|^2}{4\sigma^2}}$ where σ is the standard deviation (also called blur parameter) of the kernel which defines how large the region of influence of a particular point is, the larger the σ the bigger the region. The benefit of using kernels comes from the Moore–Aronszajn theorem [21] which states that a kernel such as the one defined above has an implicit

feature representation given by:

$$k(x_I^i, x_J^j) = \langle x_I^i, x_J^j \rangle = \varphi(x_I^i)^T \varphi(x_J^j) \quad (2)$$

As it is possible to rearrange equation 1 to be only expressed in terms of dot products, we can simply use $k(x_I^i, x_J^j)$ instead of calculating the actual dot product after the change of basis to \mathcal{H} . The modified MMD equation that makes use of kernels can be defined as:

$$MMD^2(X_I, X_J) = \mathbb{E}_{X_I}[k(x_I^i, x_I^i)] + \mathbb{E}_{X_J}[k(x_J^j, x_J^j)] - 2\mathbb{E}_{X_I, X_J}[k(x_I^i, x_J^j)] \quad (3)$$

Equation 3 provide a way of calculating pairwise discrepancy between any two WSIs with the same patch level representation. With the ability to calculate pairwise discrepancies between two WSIs, we can generate an $N \times N$ discrepancy matrix, D^{MMD} , which stores all discrepancies between all pairwise combinations of WSIs such that $D_{I,J}^{MMD} = MMD^2(X_I, X_J)$ for $I, J = 1 \dots N$. This means the relationship of the entire dataset can be represented by a single matrix that can fit into memory. We can get a corresponding symmetric and positive semi-definite Mercer kernel matrix, $K_\gamma \in R^{N \times N}$, from this discrepancy matrix as follows:

$$K_\gamma = e^{-\gamma D^{MMD}} \quad (4)$$

where $K_{I,J} = \langle X_I, X_J \rangle = \psi(X_I)^T \psi(X_J) \in [0, 1]$ is the similarity between X_I and X_J under an implicit feature representation ψ and $\gamma \in \mathbb{R}_{\geq 0}$ is a model hyper-parameter.

2.4. Implementation & Applications of MMD Kernels

To generate the MMD kernel we used GeomLoss as it provides us with an efficient GPU implementation to calculate MMD [20] and can generate the entire discrepancy matrix in around 5 hours. The implementation is available at: <https://github.com/engrodawood/Hist-MMD>. We can use this WSI level kernel for any predictive or analytical task. The development of WSI kernels enables using the vast and powerful repertoire of kernel methods in computational pathology ranging from agglomerative clustering, support vector clustering [22], kernel principal component analysis, support vector regression, classification and survival analysis. In addition it opens up the possibility of using multiple kernel learning and cross-domain feature integration. As discussed below, we demonstrate the effectiveness of the proposed approach for TP53 mutation status prediction and survival analysis.

2.5. TP53 Mutation prediction and evaluation

We assessed the significance of the proposed MMD kernel in predicting the mutation status of TP53. To do this we used a

binary class Support Vector Machine (SVM) with its pre-computed kernel set to the MMD kernel defined above [23].

For performance evaluation we used the same methodology used by the baseline method [1] which randomly partitioned the filtered TCGA slides into training, validation (used for model selection) and test sets based on 70%, 15%, and 15% ratios, respectively. The models performance was measured by using AUC-ROC. The 95% Confidence Interval (CI) of each AUC-ROC score is calculated by 1000 bootstrap runs to estimate the uncertainty of the result. The hyperparameters chosen for the kernel (through validation) are $\gamma = 4, \sigma = 10$ whereas the cost control parameter of the SVM was set as $C = 10,000$.

2.6. Kernelized Survival Analysis

We have used kernel survival SVMs (KSSVM) [24] to demonstrate the effectiveness of the proposed approach for survival analysis. KSSVM predicts a risk score $f(X)$ for a given test patient using their WSI representation X after being trained over a training dataset $\{(X_I, T_I, \delta_I) | i = 1 \dots N_{train}\}$ which consists of tuples of a patient’s whole slide image X_I , their disease-specific survival time T_i and a binary event indicator variable $\delta_I \in \{0, 1\}$ which shows whether the patient has passed away from breast cancer within a censoring time $T_{censor} = 10$ years or not. The implementation of KSSVM allows use of an implicit or pre-computed kernel and has a single hyperparameter α which controls the loss penalty term in its objective function (see [24] for further details). We chose, $\alpha = 0.125$ and γ is set to simply the multiplicative inverse of the median of D^{MMD} to reduce the number of tunable hyper-parameters as the effect of γ is also modulated by the blur parameter of the patch level kernel.

For performance evaluation we used out-of-bag bootstrapping to generate the training and test datasets $R = 50$ times stratified with respect to the survival event. In each run we generate the risk scores of the test examples which are then used to compute C-index as our primary performance metric. C-index measures the degree of concordance between relative prediction scores of test patients and their actual survival times. In line with AUC-ROC, the C-index ranges from 0.0 (inverted ranking of survival scores) to 0.5 (no concordance between predicted scores and actual survival times) to 1.0 (perfect concordance between prediction scores and actual survival times). For a representative bootstrap run, we show the Kaplan-Meier survival curves of the high risk and low risk patient groups using an optimal threshold computed from the training set scores along with p-value of the log-rank test across all bootstrap runs as: $2 \times median(\{p_r | r = 1 \dots R\})$.

3. RESULTS

The key results using the proposed approach are given in Fig. 1 with the entire TCGA-BRCA kernel matrix and the

Method	Concordance Index \pm std dev
MMD Kernel (Proposed)	0.672 ± 0.051
MCAT [12]	0.580 ± 0.069
Attention MIL [2]	0.564 ± 0.050
DeepAttnMISL [13]	0.524 ± 0.043

Table 1: Comparison of C-index \pm std dev.

resulting hierarchical clustering showing clear similarities between WSIs grouped into the same cluster.

For TP53 point mutation prediction, we obtained an AUC-ROC (with 95% CI) of 0.7958 (0.6966-0.8863) in contrast to 0.7288 (0.6205-0.8278) by the baseline study [1].

For survival prediction as can be seen in the Table 1 we achieved a much higher C-Index \pm standard deviation of 0.67 ± 0.05 in contrast to the current state of the art of 0.580 ± 0.069 [12]. Additionally, the proposed approach also allows risk stratification with Kaplan-Meier analysis with a statistically significant bootstrap p-value of the log-rank test at 0.044 as can be seen in Fig 1d.

These results show the promise of the proposed kernel as we have achieved a significantly better performance than the state of the art for two clinically important problems.

4. CONCLUSIONS AND FUTURE WORK

Beyond the application to the specific machine learning problems discussed above, the proposed MMD WSI-kernels can be used for other predictive and prognostic modelling tasks. We believe that this work will open the avenue for further research on the use of previously under-explored yet powerful kernel methods for WSI in computational pathology. In the future we will investigate the usefulness of this method on independent datasets. Second, though kernel methods provides WSI-level similarity metrics, interpretability and explainability angle of these results need to be investigated. Finally, we will also investigate the integration of multi-modal and cross-domain kernels. Our goal in this paper has been to showcase a new approach to predictive modelling of WSIs. It is important to note that clinical adoption of this and other such predictive methods warrants larger-scale testing as well as careful consideration of any confounding factors.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research was conducted retrospectively using open access human subject data from TCGA [14, 15]. Ethical approval was not required by TCGA/CBioPortal license.

6. ACKNOWLEDGMENTS

FM acknowledges funding from EPSRC EP/W02909X/1 and PathLAKE consortium. MD and FM report research funding from GlaxoSmithKline outside the submitted work.

7. REFERENCES

- [1] H. Qu et al., “Genetic mutation and biological pathway prediction based on whole slide images in breast carcinoma using deep learning,” *npj Precision Oncology*, vol. 5, no. 1, pp. 1–11, Sept. 2021, Number: 1 Publisher: Nature Publishing Group.
- [2] J. Yao et al., “Deep multi-instance learning for survival prediction from whole slide images,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I*, Berlin, Heidelberg, 2019, p. 496–504, Springer-Verlag.
- [3] G. Campanella et al., “Clinical-grade computational pathology using weakly supervised deep learning on whole slide images,” *Nature medicine*, vol. 25, no. 8, pp. 1301–1309, 2019.
- [4] J. Yan et al., “Hierarchical attention guided framework for multi-resolution collaborative whole slide image segmentation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2021, pp. 153–163.
- [5] W. Lu et al., “Slidegraph+: Whole slide image level graphs to predict her2 status in breast cancer,” *Medical Image Analysis*, p. 102486, 2022.
- [6] A. Gretton et al., “A kernel two-sample test,” *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012.
- [7] H. Sung et al., “Global cancer statistics 2020: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries,” *CA: A Cancer Journal for Clinicians*, vol. 71, no. 3, pp. 209–249.
- [8] P. May et al., “Twenty years of p53 research: structural and functional aspects of the p53 protein,” *Oncogene*, vol. 18, no. 53, pp. 7621–7636, Dec. 1999, Number: 53 Publisher: Nature Publishing Group.
- [9] Anne-Lise Børresen-Dale, “Tp53 and breast cancer,” *Human Mutation*, vol. 21, no. 3, pp. 292–300, 2003.
- [10] S. Heymann et al., “Radio-induced malignancies after breast cancer postoperative radiotherapy in patients with Li-Fraumeni syndrome,” *Radiation Oncology*, vol. 5, no. 1, pp. 104, Nov. 2010.
- [11] K. Suresh et al., “Survival prediction models: an introduction to discrete-time modeling,” *BMC medical research methodology*, vol. 22, no. 1, pp. 1–18, 2022.
- [12] R. J. Chen et al., “Multimodal co-attention transformer for survival prediction in gigapixel whole slide images,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4015–4025.
- [13] J. Yao et al., “Whole slide images based cancer survival prediction using attention guided deep multiple instance learning networks,” *Medical Image Analysis*, vol. 65, pp. 101789, 2020.
- [14] K. A. Hoadley et al., “Cell-of-origin patterns dominate the molecular classification of 10,000 tumors from 33 types of cancer,” *Cell*, vol. 173, no. 2, pp. 291–304, 2018.
- [15] The Cancer Genome Atlas Network, “Comprehensive molecular portraits of human breast tumours,” *Nature*, vol. 490, no. 7418, pp. 61–70, Oct 2012.
- [16] E. Cerami et al., “The cbio cancer genomics portal: an open platform for exploring multidimensional cancer genomics data,” *Cancer discovery*, vol. 2, no. 5, pp. 401–404, 2012.
- [17] J. Gao et al., “Integrative analysis of complex cancer genomics and clinical profiles using the cbiportal,” vol. 6, no. 269, pp. p11–p11, 2013.
- [18] J. Liu et al., “An integrated tcga pan-cancer clinical data resource to drive high-quality survival outcome analytics,” *Cell*, vol. 173, no. 2, pp. 400–416, 2018.
- [19] X. Zhang et al., “Shufflenet: An extremely efficient convolutional neural network for mobile devices,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6848–6856.
- [20] J. Feydy et al., “Interpolating between optimal transport and mmd using sinkhorn divergences,” in *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019, pp. 2681–2690.
- [21] G. Wahba, “An introduction to reproducing kernel hilbert spaces and why they are so useful,” in *Proceedings of the 13th IFAC Symposium on System Identification (SYSID 2003)*, 2003.
- [22] A. Ben-Hur et al., “Support vector clustering,” *Journal of machine learning research*, vol. 2, no. Dec, pp. 125–137, 2001.
- [23] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [24] S. Pölsterl et al., “Fast training of support vector machines for survival analysis,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2015, pp. 243–259.