
Extending Fair Null-Space Projections for Continuous Attributes to Kernel Methods

Anonymous Authors¹

Abstract

With the on-going integration of machine learning systems into the everyday social life of millions the notion of fairness becomes an ever increasing priority in their development. Fairness notions commonly rely on protected attributes to assess potential biases. Here, the majority of literature focuses on discrete setups regarding both target and protected attributes. The literature on continuous attributes especially in conjunction with regression – we refer to this as *continuous fairness* – is scarce. A common strategy is iterative null-space projection which as of now has only been explored for linear models or embeddings such as obtained by a non-linear encoder. We improve on this by extending this to kernel induced feature spaces by means of the “empirical feature space”. We theoretically derive this as a direct transformation of the kernel matrix yielding a model and fairness-score agnostic method applicable to continuous protected attributes. We demonstrate that our novel approach in conjunction with Support Vector Regression (SVR) provides competitive or improved performance across multiple datasets in comparisons to other contemporary methods.

1. Introduction

“The development, deployment and use of AI systems must be fair” is a central statement in the Ethics Guidelines for Trustworthy AI of the EU Commission (2019, p. 12). Similarly fair treatment is in many areas governed by anti-discrimination laws such as Article 21 in the charter of fundamental rights of the European Union (2012) which among others include protection against discrimination based on “age”, “sex” and “race”. Evidently, in recent years significant attention has been drawn to this issue not only in the

ML community itself but also by legislators. This led to a series of developments to ensure the fairness of ML systems.

ML systems typically make use of all data provided to them. As a result, they can inherit biases – such as those stemming from ageism or racism – and may learn to reproduce these patterns. Fair ML aims to address this issue by attempting to mitigate or remove such biases. Most commonly, this is done by explicitly marking certain features as protected attributes – for example, “age”, “sex” or “race” – which should not influence the learned decision rules, such as determining “creditworthiness” or the “crime probability”. The challenge, however, lies in the fact that even if these attributes are not directly included in the training data, they can still be inferred indirectly through hidden correlations, allowing the system to pick up on them unintentionally.

Most often, the focus is on categorical protected attributes (such as “race”) in a classification setting. Whilst a large body of literature assumes this setting, an attribute such as “age” is not naturally categorical and other attributes such as “race” are commonly found as a population percentage in a given community which is then often categorized as pre-processing. Further, the more general regression setting is also not as frequently considered thus rendering the available options to ensure the by law required non-discrimination guarantees slim. We will refer to the setup in which the targets as well as the protected attributes are continuous rather than discrete as “*continuous fairness*”.

A common strategy in fairness is the idea of “null-space projection”: first a direction in space predictive of the protected is determined and then the data projected so that this direction becomes uninformative. An iterative process that each time removes different pieces of information. However, this idea is thus far limited to either linear methods or the application to non-linear embeddings. This work significantly extends the approach to kernel induced feature spaces, specifically our contributions are the following:

- We derive a method to perform a direct transformation of the kernel matrix that corresponds to a null-space projection in the empirical feature space. We show that the theoretical properties of the kernel matrix are retained and demonstrate out-of-sample extension.

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

- We illustrate that the transformed kernel is model agnostic: most notably giving rise to a method for fair Support Vector Regression.
- We empirically obtain competitive or improved performance against other contemporary methods on a range of different datasets and fairness scores.

This paper is structured as follows: first we introduce the field of fair machine learning by discussing a selection of related works (Section 2) with a focus on social implications in Section 2.3 before presenting the extension of null-space projections to kernel methods (Section 3). In Section 4 we empirically demonstrate competitive performance on real-world datasets, afterwards discussing our approach, its limitations and directions for future work (Section 5).

2. Related Work

We first discuss the general setting of fair ML before diving into specific fairness approaches related to subspace projections and kernel methods.

2.1. Fairness in Machine Learning

Two main stream characterization of fairness in machine learning can be distinguished: *individual* and *group* fairness. The former aims for similar predictions for similar individuals the latter looks at the group specific predictions given a so-called protected attribute (such as “race” or “age”). For an extensive discussion see Binns (2020). The most common measure for fairness are *Demographic/Statistical Parity* and *Equalized Odds*. The former aims at an equal probability to receive an advantageous decision/label (e.g. credit worthiness) whilst the latter also takes the true label into account. It is not possible to directly compute these for the continuous fairness setting and hence adequate adaptations are required which we discuss next.

2.2. Continuous Fairness Measures

In recent years, a set of novel approaches to compute fairness measures for continuous fairness have been introduced – most aim to measure the statistical dependence between the model prediction and the protected attribute. Mary et al. (2019) introduce a fairness measure based on the Hirschfeld–Gebelein–Rényi (HGR) maximum correlation and a method for its estimation based on kernel density estimations. Concurrently, two generalizations of demographic parity were introduced: Giuliani et al. (2023) present (inspired by properties of the HGR) a family of Generalized Disparate Impact indicators (GeDI) and Jiang et al. (2022) put forward “Generalized Demographic Parity” (GDP) by using a kernel density estimation to weight local and global prediction averages.

A different approach is taken by Narasimhan et al. (2020) by deriving pairwise fairness (PF) a measure based on pairwise comparisons for demographic parity and equalized odds.

It has been mathematically shown (Kleinberg et al., 2017) that the respective objectives of different fairness measures are mutually exclusive. It therefore remains up to the designer of the machine learning system which measure to employ during training. This is aggravated in the continuous fairness setting as computational tractability becomes an issue: complexity usually ranges from quadratic to cubic.

One strategy is to include one of the fairness measures (or a suitable surrogate) into the optimization problem as a penalty - for example Mary et al. (2019) use a neural network penalized by an upper bound of the HGR (denoted “NN-HGR”), Kong et al. (2025) penalize with a new metric whilst Lai & Guan (2025) rely on an adversarial technique.

2.3. Continuous Fairness: Social Implications

This section aims to explore how the wider body of literature including social sciences deal with the issue of continuous fairness and how it is recognized therein.

Stypinska (2023) discuss the general topic of ageism related to AI systems and how this on the technical level can be attributed to under-representation or cutoffs of certain age groups whilst often being neglected in the fairness discourse. Chu et al. (2023) perform an extensive survey to explore the leading sources of age-related biases in AI systems concluding among other things that “the second phenomenon is the misinterpretation of age as a category [...]” (p. 12). The effect of categorizing “age” has been recognized by Hort & Sarro (2022) yet the focus rather remains on finding suitable categories than treating it as a continuous quantity.

Regarding the attribute “race”, many works in the social and related sciences commonly work with it as a population percentage, for example when analyzing the connection to fatal police shootings (Zare et al., 2025) or traffic citations (Xu et al., 2024). Note that these works rely on data from the “American Community Survey” which also inspires two of the three datasets used in this work for empirical evaluation.

This demonstrates a clear disparity between how the attributes are dealt with by experts in practice and how they are (often for algorithmic convenience) treated within the ML community.

2.4. Sub-space Projections for Fairness

Subspace projections comprise methods aimed at identifying subspaces with specific properties usually to ensure subsequently applied algorithms satisfy certain guarantees. This a common approach to ensure workspace constraints in robotics (Dietrich et al., 2015).

In the domain of fairness Pérez-Suay et al. (2017) also introduce a fair dimensionality reduction technique based on the “Hilbert-Schmidt independence criterion” (HSIC) (Gretton et al., 2007).

Tan et al. (2020) take the viewpoint of empirical risk minimization trying to find a subspace of the hypothesis space that fulfills predictive but also fairness properties for which they derive an approach for Gaussian Processes.

Ravfogel et al. (2020) on the other hand suggests to first find a direction in space deemed relevant for predicting the protected attribute - in their case the weights of a linear classifier¹ whose null-space is then used for the projection of either the data or their embedding in a neural network.

This work generalizes the work of Ravfogel et al. (2020) by considering the regression setting for kernel methods i.e. we provide means to work with null-space projections for the non-linear features induced by kernels including cases where the corresponding feature space is infinite dimensional which is a novel contribution. As our approach focuses on kernel methods we discuss previous work on fairness with kernels in the next section.

2.5. Fairness in Kernel Methods

There is some work regarding the continuous fairness setting with kernels in the literature: Pérez-Suay et al. (2017) also develop a fair approach for Kernel Ridge Regression (we refer to this as “KRR-FKL”) using the HSIC which is extended to Gaussian Processes but without out-of-sample extension (Perez-Suay et al., 2023). Oneto et al. (2020) introduce constraints to the KRR optimization to develop an alternative approach.

In contrast, even though there exist a variety of different approaches for fair Support Vector Machines (SVMs) (Olfat & Aswani, 2018; Donini et al., 2018; Zafar et al., 2019; Carrizosa et al., 2023) they are neither extendable to regression nor to continuous protected attributes – at least without more intricate changes as they mostly rely on modification of the original optimization problem to incorporate different measures of fairness as constraints or penalties.

The method introduced next is model agnostic and applicable for fairness as long as the approach relies on kernels.

3. Method

In contrast to other approaches, we do not directly optimize for a specific fairness-score. Rather, we aim to remove information that is required to predict the protected attribute. For

¹They also mention that this can be extended to the regression setting, however best to the authors knowledge this has not yet been thoroughly investigated in the community.

this we extend the idea of null-space projections presented by Ravfogel et al. (2020) to the setting of kernel methods.

3.1. Motivation: Feature Space

At the heart of most kernel methods lies the kernel trick: many classical algorithms can ultimately be expressed solely in terms of inner products between data-points. Kernels exploit this by enabling efficient computation of inner products in potentially infinite-dimensional spaces. Formally, they are equivalent to performing an implicit pre-processing step in a (possibly infinite-dimensional) feature space. A kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a map that factors through some inner product, i.e.

$$k(\mathbf{x}, \mathbf{x}') = \langle \Phi(\mathbf{x}), \Phi(\mathbf{x}') \rangle \tag{1}$$

where Φ is the feature map to some Hilbert space (the corresponding Hilbert space is referred to as feature space, the original space as input space). The most well known kernel is the rbf-kernel given by

$$k_{rbf}(\mathbf{x}, \mathbf{x}') := \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|^2) \tag{2}$$

with parameter γ . This is also the one considered in this work and it does correspond to an infinite-dimensional feature space (Shawe-Taylor, 2004).

The naive extension of null-space projections to the respective feature spaces is only applicable if the dimension is finite and only tractable if quite small – as it is impossible to directly work with data living in an infinite dimensional space. However there are two distinct ways this issue can be tackled: we usually only require the inner product between two projected points which can be implicitly expressed as a kernel recurrence or by using the so-called *empirical feature space* which is the approach taken in the following.

3.2. Fairness in the Empirical Feature Space

Given n samples let $\mathbf{K} \in \mathbb{R}^{n \times n}$ denote the kernel matrix induced by a kernel $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ applied to the d -dimensional data stored in $\mathbf{X} \in \mathbb{R}^{n \times d}$. Our goal is to remove information about the protected attribute from \mathbf{K} whilst keeping its required properties. As \mathbf{K} is symmetric positive-semidefinite it decomposes as

$$\mathbf{K} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\top = \mathbf{G}\mathbf{G}^\top \tag{3}$$

using the eigendecomposition of \mathbf{K} and defining $\mathbf{G} := \mathbf{Q}\mathbf{\Lambda}^{1/2}$ where $\mathbf{\Lambda}^{1/2} := \text{diag}(\sqrt{\lambda_1}, \dots, \sqrt{\lambda_n})$. Note that the columns of \mathbf{Q} correspond to the eigenvectors of \mathbf{K} . We view $\mathbf{G} = (\mathbf{g}_1, \dots, \mathbf{g}_n)^\top \in \mathbb{R}^{n \times n}$ as an abstract data representation for the linear kernel, i.e., $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \mathbf{g}_i, \mathbf{g}_j \rangle$ and \mathbf{g}_i being a column vector. This means that we can write $\mathbf{G}\mathbf{G}^\top = \mathbf{K}$. This is called the empirical feature space which provides an alternative view on the kernel trick and is

defined as $\mathbf{Y} := \mathbf{K}\mathbf{Q}\mathbf{A}^{-1/2}$ (Scholkopf et al., 1999; Xiong et al., 2005; Kwak, 2013). The connection to Equation (3) can be seen by observing that

$$\mathbf{K}\mathbf{Q}\mathbf{A}^{-1/2} = \mathbf{Q}\mathbf{A}\mathbf{Q}^\top\mathbf{Q}\mathbf{A}^{-1/2} = \mathbf{Q}\mathbf{A}^{1/2} = \mathbf{G}. \quad (4)$$

This *empirical feature space* retains the geometric structures of the finite subspace spanned by the training data within the feature space (it is isometric) and corresponds to the subspace that a kernel method relies on for training (Scholkopf et al., 1999). As $\mathbf{g}_1, \dots, \mathbf{g}_n$ is a finite dimensional representation we can directly apply subspace projections and apply the linear kernel afterwards. This is now done by null-space projections as presented by Ravfogel et al. (2020) for the case of regression. That is first we find a weighting that contains information for predicting the protected attribute \mathbf{p}

$$\min_{\mathbf{w}} \|\mathbf{p} - \mathbf{G}\mathbf{w}\|^2 + \tilde{\alpha}\|\mathbf{w}\|^2 \quad (5)$$

by a default ridge-regression approach. The solution to this is straight-forward and well known as

$$\mathbf{w} = \mathbf{p}^\top(\mathbf{G}\mathbf{G}^\top + \tilde{\alpha}\text{Id})^{-1}\mathbf{G}. \quad (6)$$

Note that there are two equivalent ways of writing Equation (6), this version is required for the kernel trick. Similarly, we require this so the projection can later be expressed in terms of \mathbf{K} . Then we remove the information from \mathbf{G} by enforcing

$$\mathbf{G}\mathbf{w} \stackrel{!}{=} \mathbf{0} \quad (7)$$

with the null-space projection $\mathbf{P}^{\mathbf{G}}$ as follows:

$$\mathbf{G}_{(1)} = \mathbf{G}_{(0)}\mathbf{P}_{(0)}^{\mathbf{G}} \quad (8)$$

where we define $\mathbf{P}_{(1)}^{\mathbf{G}} := (\text{Id} - \mathbf{w}(\mathbf{w}^\top\mathbf{w})^{-1}\mathbf{w}^\top)$ as this enforces Equation (7) and where the scalar $(\mathbf{w}^\top\mathbf{w})^{-1}$ ensures normalization. Of course $\mathbf{G}_{(1)}$ might still contain information about the protected attribute so this process can be iterated by finding a new \mathbf{w} and chaining the projection matrices:

$$\mathbf{G}_{(m)} = \mathbf{G}_{(0)} \prod_{i=0}^{m-1} \mathbf{P}_{(i)}^{\mathbf{G}}. \quad (9)$$

The projected kernel matrix at the m -th iteration is then retrieved by

$$\mathbf{K}_{(m)} = \mathbf{G}_{(m)}\mathbf{G}_{(m)}^\top. \quad (10)$$

which can now be used as a kernel for training a prediction task with respect to the target. Note that $\mathbf{K}_{(0)}$ corresponds to the original kernel matrix. Yet, it is not clear whether the product $\prod_{i=0}^{m-1} \mathbf{P}_{(i)}^{\mathbf{G}}$ is a valid projection matrix as in general the product of projections does not need to be a projection itself. This is ensured by the subsequent lemma following the approach taken in Ravfogel et al. (2020).

Lemma 3.1. $\prod_{i=0}^{m-1} \mathbf{P}_{(i)}^{\mathbf{G}}$ is a projection.

Proof. All proofs are found in Appendix A. \square

To perform the projection for k unseen test points they have to be first mapped to the empirical feature space (see Xiong et al. (2005); Kwak (2013)) then the projection in Equation (9) performed and the linear kernel applied. This detour can be avoided by expressing this in a single transformation directly in terms of the kernel matrix which is one of our key contributions and is given by the following theorem:

Theorem 3.2. Let $\mathbf{K}_{(m)}$ be defined as in Equation (10). Then for $m > 0$ it holds:

$$\mathbf{K}_{(m)} = \mathbf{K}_{(m-1)}\mathbf{T}^{\mathbf{K}_{(m-1)}} \quad (11)$$

where we set $\tau_{norm} := (\mathbf{w}^\top\mathbf{w})^{-1}$ and define

$$\mathbf{T}^{\mathbf{K}_{(m)}} := (\text{Id} - \mathbf{M}\mathbf{K}_{(m)})$$

$$\mathbf{M} := (\mathbf{K}_{(m)} + \tilde{\alpha}\text{Id})^{-1}\mathbf{p}\tau_{norm}\mathbf{p}^\top(\mathbf{K}_{(m)} + \tilde{\alpha}\text{Id})^{-1}.$$

Algorithm 1 provides a pseudo-code of this procedure. The following corollary emphasizes that the resulting kernel matrix *firstly* still corresponds to the inner product in some feature space and thus *secondly* ensure necessary properties such as convexity in subsequent algorithms (such as SVR).

Corollary 3.3. $\mathbf{K}_{(m)}$ in eq. (11) is positive semi-definite.

Note that this also extends to the case of multiple protected attributes (see Appendix A.2). Theorem 3.2 yields a transformation $\mathbf{T}^{\mathbf{K}_{(m-1)}}$ for a single iteration which removes the relevant information on per-row basis whilst maintaining positive semi-definiteness². Most importantly this can be directly applied to $\mathbf{K}_{test} \in \mathbb{R}^{k \times n}$. The transformations can be combined to yield a single transformation w.r.t $\mathbf{K}_{(0)}$ as

$$\mathbf{T}_m := \prod_{i=0}^{m-1} \mathbf{T}^{\mathbf{K}_{(i)}}. \quad (12)$$

so that it holds $\mathbf{K}_{(m)} = \mathbf{K}_{(0)}\mathbf{T}_m$. When we refer to iteration m of our approach we mean the transformation \mathbf{T}_m . Note that \mathbf{T}_m corresponding to the linear kernel applied to the iterated projection of the empirical feature space – \mathbf{T}_m itself is not a projection. Alternatively it is possible to look at this directly in the feature space but this is beyond the scope of this work. Note that in contrast to Pérez-Suay et al. (2017) we directly find a predictive direction as weights of a model and Tan et al. (2020) instead aims to align different subspaces of the hypothesis space.

²A kernel k is usually called valid or positive definite if the corresponding kernel matrix is positive semi-definite for all finite number of points.

Algorithm 1 Fair Kernel Decomposition

```

1: Input:  $n \times n$  Kernel matrix  $\mathbf{K}$ ,  $n \times l$  protected attribute(s)  $\mathbf{p}$ , number of iterations  $m$ , regularization parameter  $\tilde{\alpha}$ .
2:  $\mathbf{K}_0 \leftarrow \mathbf{K}; \mathbf{T}^{\mathbf{K}_0}, \mathbf{T}_0 \leftarrow \text{Id}$ 
3: for  $i \leftarrow 1$  to  $m$  do
4:    $\mathbf{B} \leftarrow (\mathbf{K}_{(i-1)} + \tilde{\alpha}\text{Id})^{-1}$  {Nystroem possible}
5:    $\tau_{norm} \leftarrow (\mathbf{p}^\top \mathbf{B} \mathbf{K}_{(i-1)} \mathbf{B} \mathbf{p})^{-1}$ 
6:    $\mathbf{M} \leftarrow \mathbf{B} \mathbf{p} \tau_{norm} \mathbf{p}^\top \mathbf{B}$ 
7:    $\mathbf{T}^{\mathbf{K}_i} \leftarrow \text{Id} - \mathbf{M} \mathbf{K}_{(i-1)}$  {Trf. from  $\mathbf{K}_{(i-1)}$  to  $\mathbf{K}_{(i)}$ }
8:    $\mathbf{K}_{(i)} \leftarrow \mathbf{K}_{(i-1)} \mathbf{T}^{\mathbf{K}_i}$ 
9:    $\mathbf{T}_{(i)} \leftarrow \mathbf{T}_{(i-1)} \mathbf{T}^{\mathbf{K}_i}$  {Trf. from  $\mathbf{K}_{(0)}$  to  $\mathbf{K}_{(i)}$ }
10: end for
11: Output: Transformed positive semi-definite kernel matrix  $\mathbf{K}_{(m)}$ , transformation matrix  $\mathbf{T}_{(m)}$ .
    
```

3.2.1. CHOICE OF ALGORITHM

Even though we rely on a ridge-regression procedure for the null-space projection this method only directly affects the kernel so it can be plugged into any subsequent algorithm relying on a kernel matrix. In this work we focus on the two canonical classical approaches: Kernel Ridge Regression (KRR) and Support Vector Regression (SVR).

3.2.2. COMPUTATIONAL COMPLEXITY AND NYSTROEM

In terms of complexity our approach requires to invert the $n \times n$ kernel matrix (usually $\mathcal{O}(n^3)$). This expensive operation can be avoided using the nystroem approximation (Drineas & Mahoney, 2005) for the inverse $(\mathbf{K}_{(m-1)} + \tilde{\alpha}\text{Id})^{-1}$ in Theorem 3.2. Further, computational complexity from matrix multiplications can be reduced by choosing an appropriate order for multiplication and by not storing the transformations $\mathbf{T}_{(i)}$ explicitly – we provide the full details in Appendix C due to space restrictions.

We evaluate the nystroem approximation in Section 4.4 – it however remains a challenge to investigate the true scalability as two of three considered fairness measures (that are not related to our method but are required to quantitatively evaluate continuous fairness) are in $\mathcal{O}(n^3)$.

4. Experimental Results

All implementations as well as all experiments in this chapter are made publicly available on Github³.

4.1. Real-world Datasets

We consider the following datasets:

³After-anonymous-review

- *Communities & Crimes*: This dataset is provided by Redmond & Baveja (2002) and widely used in the continuous fairness setup. The premise is to predict the “crime rate” based on characteristics of a given community. The protected attribute is the “percentage of black people”.
- *ACSIncome*: Ding et al. (2021) introduce novel datasets based on data by the “American Community Survey” and is distinguished by US state. The “ACSIncome” task aims to predict the income of a person based on education, work hours and other social characteristics. We consider “age” as protected attribute. The size of the data varies considerably between different states, to render computations tractable we choose the smaller dataset corresponding to “Montana”.
- *ACSTravelTime*: This is also introduced by Ding et al. (2021) with the aim of predicting the “time to commute” based on social and personal characteristics. We consider “age” as protected attribute and also choose the state “Montana”.

Unfortunately there is a lack of datasets for fairness in general and it is up to debate how sensible (see for example Bao et al. (2021)) existing and widely used datasets are – an issue exacerbated by the “continuous fairness” setting. However, the ACS data is commonly used in social fields (refer to Section 2.3). Each problem discussed above aims at possible use-cases in the real-world: city planners develop efficient traffic routing (“ACSTravelTime”), employers predict the salary of a new hire (“ACSIncome”) and the police relegates its forces to certain regions (“Crimes”). In all these cases laws govern that there should be no discrimination based on “age” and “race” present as continuous attributes.

4.1.1. METHODS

We compare our approach in combination with KRR and SVR which we term “KRR/SVR-FKD” where “FKD” stands for “Fair Kernel Decomposition”. To benchmark against another fairness approach for KRR we compare against Pérez-Suay et al. (2017) where we denote their approach as “KRR-FKL”. Additionally we benchmark against the HGR penalized neural network from Mary et al. (2019) (which we denote “NN-HGR”) as well as a dummy classifier that simply predicts the training’s targets mean.

We do not compare against Perez-Suay et al. (2023) as they require label information for test data and found the most recent approach (Kong et al., 2025) to not outperform the other baselines (see Appendix D for a comparison).

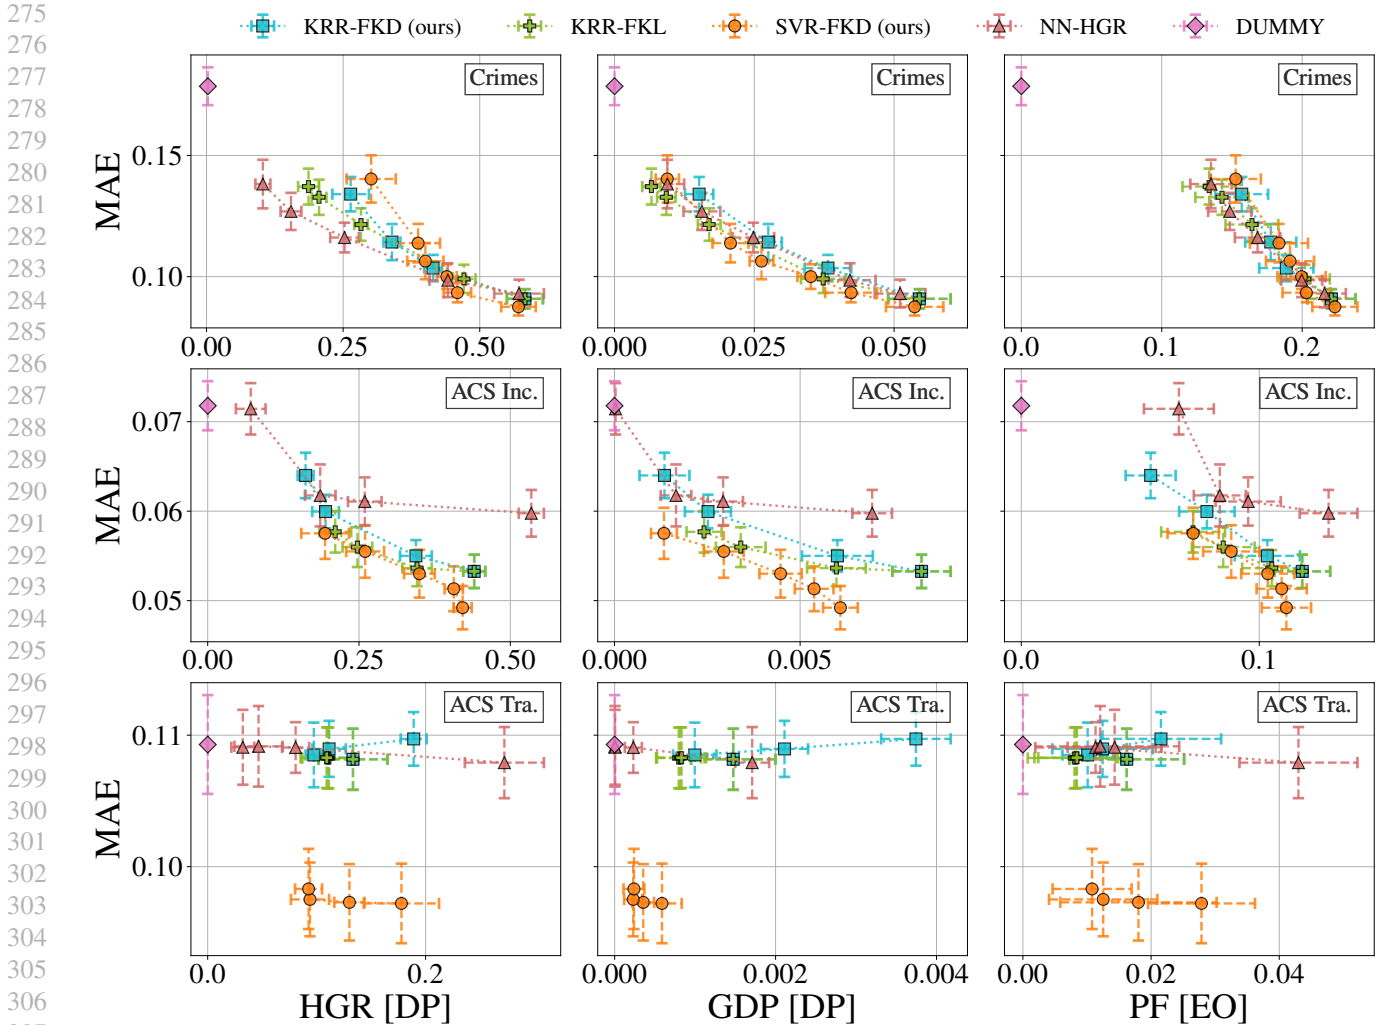


Figure 1. Pareto fronts for the MAE (y-axis, lower is better) and three different fairness measures (x-axis, lower is better) reported with empirical standard deviations (error-bars). Comparison for multiple methods (marker, color) and datasets (boxes, one per row). Number of iterations m for *-FKD, regularization μ for *-FKL and HGR penalty λ are found in Appendix B.

4.1.2. PARAMETERS

For the kernel methods we use the rbf-kernel (see Equation (2)). For every dataset we first run a grid-search on the non-fair base methods (KRR, SVR) to find their optimal default hyper-parameters. For our approach we have an additional parameter $\tilde{\alpha}$ for the KRR used inside our decomposition approach which we set heuristically (for an analysis see Section 4.3). Otherwise we run our evaluation for a range of different regularization parameters (penalty coefficient and number of iterations). Details about the used parameters and more is given in Appendix B. For “NN-HGR” we choose the architecture reported in the original contribution (Mary et al., 2019)⁴.

⁴Note that this architecture might not be optimal for every dataset but a fair comparison of completely different approaches is not straight-forward for multi-objective problems due to the com-

4.1.3. MEASURES

All reported results are obtained by a 5-fold cross validation. We report the mean absolute error (MAE) for target evaluation and use the following measures for fairness (see Section 2.1 for the corresponding explanation and references): to evaluate for demographic parity we choose HGR (denoted HGR [DP]) and GDP (denoted GDP [DP]) whereas for equalized odds⁵ we use pairwise fairness (denoted PF [EO]). The empirical standard deviation across all runs is provided for every measure and shown as error-bars.

plexity of hyperparameter optimization. This method is anyways included to embed our approach into the wider body of literature.

⁵For pairwise fairness this is identical to equal opportunity.

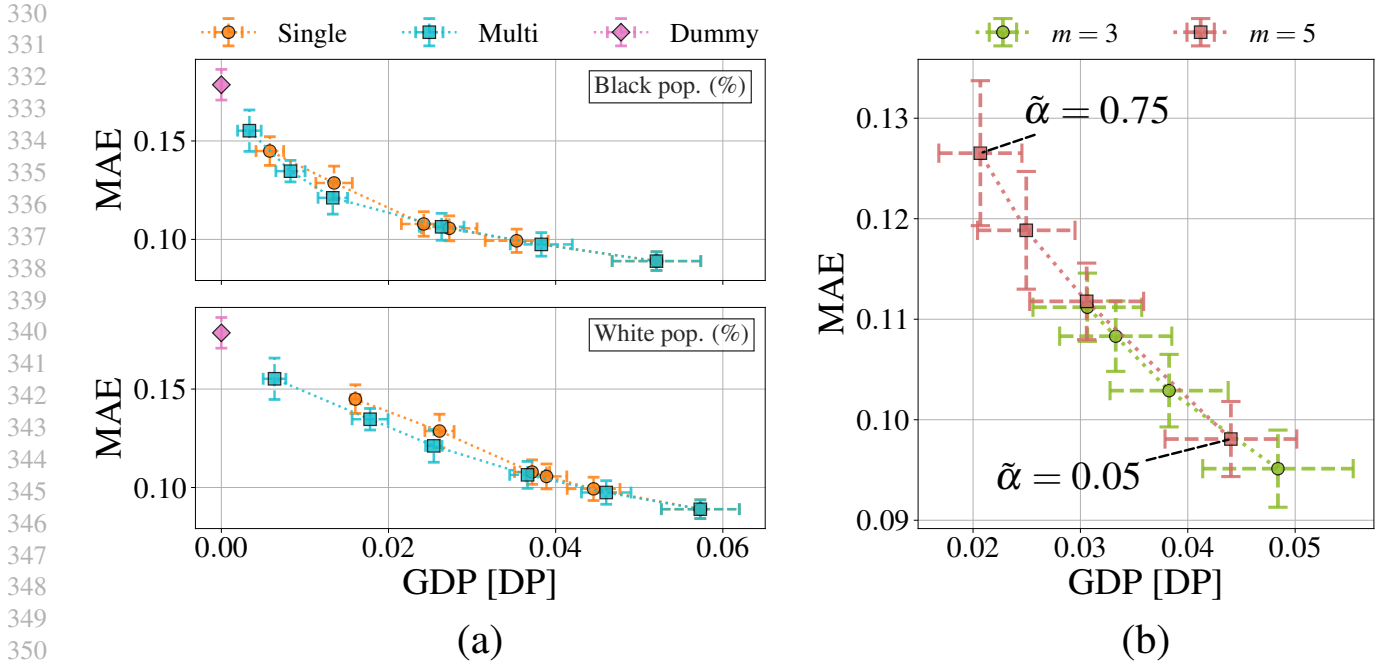


Figure 2. Other experiments on the “Crimes” dataset. (a): Performance (MAE, y-axis, lower is better) of our “SVR-FKD” considering both the white and black population percentage (pop. perc.) as protected (multi, squares, cyan, $m \in \{0, 5, 12, 25, 35, 45\}$) and only the black pop. perc. (single, circles, orange, $m \in \{0, 10, 35, 50, 70, 85\}$). GDP fairness (x-axis, lower is better) with respect to the black pop. perc. (top) and to the white pop. perc. (bottom) (b): Influence of $\tilde{\alpha}$ on the performance (MAE, y-axis, lower is better) of our “KRR-FKD” for fixed parameters and iterations $m \in \{3, 5\}$ (color and marker). GDP fairness (x-axis, lower is better) w.r.t the black pop. perc.

4.1.4. RESULTS

Evaluations concerning fair machine learning require to consider a trade-off between predictive performance and fairness – hence we are interested in the corresponding (approximate) pareto fronts which are visualized in Figure 1. Note that each fairness measure is taken at the same time for the same model i.e. every row evaluates the same models and same predictions but with different fairness measures.

On the *Crimes* dataset differences can be observed for the HGR measure where “NN-HGR” and “KRR-FKL” outperform our approach in case of strong regularization despite that this margin is interestingly much closer for the other measures and our “SVR-FKD” outperforms the other approaches slightly for weak regularization being more pronounced for the GDP score.

On the *ACSIncome* dataset our “SVR-FKD” generally outperforms all other approaches particular pronounced for GDP, whereas our “KRR-FKD” is outperformed by “KRR-FKL” both in turn providing better trade-offs than “NN-HGR” even for the HGR score.

On the *ACSTravelTime* dataset we observe that most methods – even if unregularized – perform equally to the dummy regressor. This is not the case for “SVR-FKD” that also achieves considerable improvements in terms of fairness.

In conclusion we can see quite heterogeneous results – emphasizing that to achieve fairness the right method has to be chosen for the right task-at-hand. Our newly introduced methods show competitive performance across multiple measures and datasets with best performance for “SVR-FKD” thus providing an useful new method for fair ML.

4.2. Multiple Protected Attributes

Next to predictive performance we so-far optimized for fairness with respect to a single attribute. Of course it is possible we either miss to ensure fairness with respect to another attribute or maybe even introduce unwanted biases by discriminating against other groups.

4.2.1. EXPERIMENTAL SETUP

To analyze how our approach compares when considering multiple protected attributes we again utilize the “Crimes” dataset. Next to the “black population percentage” (pop. perc.) it also includes the “white population percentage”. We compare the performance of our “SVR-FKD” once being trained to protect both attributes (“multi”) and once only for the “black population percentage” (“single”) and then measuring fairness with respect to each attribute. Otherwise the same experimental setup is used as previously.

4.2.2. RESULT

Figure 2 (a) shows how in both cases fairness improves for both groups which is sensible as there is a clear connection between the attributes. For the black pop. perc. we observe very similar fairness-scores but with slight differences for high regularization – this effect is more pronounced for the white pop. perc. where we do observe more noticeable differences as the multi-protected approach achieves improved trade offs – even though the margin remains slim. The multi-protected setting is of interest for further investigation with wider availability of appropriate datasets in the future.

 4.3. Influence of $\tilde{\alpha}$

A key parameter in our approach is the regularization $\tilde{\alpha}$ for the ridge-regression that computes the weight used for the null-space projection. Usually $\tilde{\alpha}$ controls for overfitting, however we might benefit from only removing very specific information about the protected attribute – rendering the interpretation more difficult.

We analyze this by measuring the performance of “KRR-FKD” for a fixed number of iterations on the “Crimes” dataset whilst varying $\tilde{\alpha}$ – this is shown in Figure 2 (b). The parameter plays a crucial role: the higher $\tilde{\alpha}$ the larger the change to predictive performance i.e. smaller values correspond to a more fine-grained removal of information.

This suggests to use a smaller value and control the degree of regularization by choosing an appropriate m – the precise choice however as always remains dataset specific.

4.4. Nystroem Approximation

To address the scalability issues we test the performance difference if we rely on the nystroem inverse approximation. We use the same settings as before, focus on the “Crimes” dataset and show the performance for different levels of approximation by using a certain percentage of the original components for the approximation.

Figure 3 clearly demonstrates qualitatively close results even if some differences persist the overall trend is identical. It is of great interest to investigate the limits of this scalability in future – however the evaluation of fairness measures remains a challenge for large datasets as many measures suffer from cubic complexity.

5. Conclusion

This work introduces an approach that allows the extension of null-space projections for fairness to the more general case of kernel methods for regression with continuous protected attributes. We term this setting “continuous fairness” which is often neglected within the ML community (Sec-

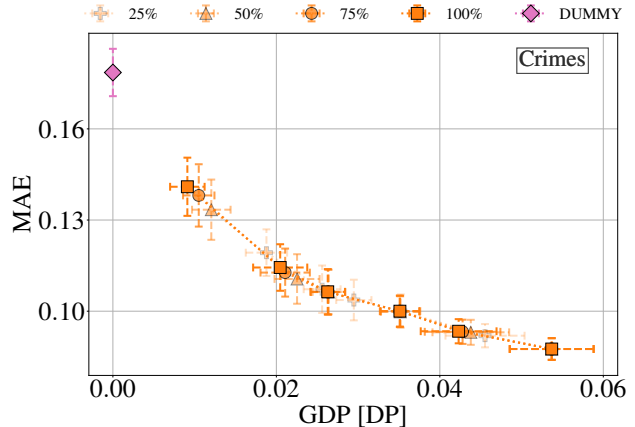


Figure 3. Nystroem approximation with different percentages of components for “SVR-FKD”. “Crimes” dataset with the same experimental setting as before.

tion 2.3). Our approach is model and task agnostic as it directly depends on the kernel matrix allowing for a natural extension to unseen data.

We empirically demonstrate that our approach is suited for the application of different regression strategies using the transformed kernel matrix. This yields competitive and in case of our newly introduced approach for fair Support Vector Regression with continuous protected attributes (“SVR-FKD”) improved performance.

5.1. Limitations and Future Work

Some limitations are inherent to our approach, such as the iterative and kernel based nature: for m iterations and n samples we require $\mathcal{O}(m \cdot n^3)$ operations. We empirically found that this can be significantly reduced by the nystroem method for matrix inversion without losing performance: however due to the limited availability of scalable continuous fairness measures, true scalability needs a more intricate investigation in the future. On this end, it is of interest to analyze the possibility of directly performing the nullspace projections in the reduced nystroem representation.

Future work might focus on analyzing the performance on other *kernel methods* we did not investigate in this work such as Gaussian Processes.

Another interesting extension and modification of our approach would be to other *tasks* such as classification with a continuous protected attribute, regression with a discrete protected attribute or in other domains such as privacy.

A future analysis of the alternative view point of direct projection in the feature space is another direction of interest.

Impact Statement

This paper presents work whose goal is to advance the field of Fair Machine Learning – we do not argue in favor or against the use of such systems in domains of social impact but reality is that ML systems do strongly affect society. Further, as with most fairness approaches it is possible to misuse the introduced algorithm – for example to wrongfully remove the contributions of marginalized groups.

References

- Bao, M., Zhou, A., Zottola, S., Brubach, B., Brubach, B., Desmarais, S., Horowitz, A., Lum, K., and Venkatasubramanian, S. It’s COMPASlicated: The Messy Relationship between RAI Datasets and Algorithmic Fairness Benchmarks. In Vanschoren, J. and Yeung, S. (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.
- Binns, R. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* ’20*, pp. 514–524, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 978-1-4503-6936-7.
- Carrizosa, E., Halskov, T., and Romero Morales, D. Wasserstein SVM: Support vector machines made fair, June 2023.
- Chu, C. H., Donato-Woodger, S., Khan, S. S., Nyrup, R., Leslie, K., Lyn, A., Shi, T., Bianchi, A., Rahimi, S. A., and Grenier, A. Age-related bias and artificial intelligence: A scoping review. *Humanities and Social Sciences Communications*, 10(1):510, August 2023. ISSN 2662-9992.
- Commission, E. *Ethics Guidelines for Trustworthy AI*. Publications Office, 2019.
- Dietrich, A., Ott, C., and Albu-Schäffer, A. An overview of null space projections for redundant, torque-controlled robots. *The International Journal of Robotics Research*, 34(11):1385–1400, 2015.
- Ding, F., Hardt, M., Miller, J., and Schmidt, L. Retiring Adult: New Datasets for Fair Machine Learning. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P. S., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 6478–6490. Curran Associates, Inc., 2021.
- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J., and Pontil, M. Empirical risk minimization under fairness constraints. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, pp. 2796–2806, Montréal, Canada and Red Hook, NY, USA, 2018. Curran Associates Inc.
- Drineas, P. and Mahoney, M. W. On the Nystrom Method for Approximating a Gram Matrix for Improved Kernel-Based Learning. *Journal of Machine Learning Research*, 6(72):2153–2175, 2005.
- Giuliani, L., Misino, E., and Lombardi, M. Generalized Disparate Impact for Configurable Fairness Solutions in ML. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 11443–11458. PMLR, July 2023.
- Gretton, A., Fukumizu, K., Teo, C., Song, L., Schölkopf, B., and Smola, A. A Kernel Statistical Test of Independence. In Platt, J., Koller, D., Singer, Y., and Roweis, S. (eds.), *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- Hort, M. and Sarro, F. Privileged and Unprivileged Groups: An Empirical Study on the Impact of the Age Attribute on Fairness. In *2022 IEEE/ACM International Workshop on Equitable Data & Technology (FairWare)*, pp. 17–24, 2022.
- Jiang, Z., Han, X., Fan, C., Yang, F., Mostafavi, A., and Hu, X. Generalized Demographic Parity for Group Fairness. In *International Conference on Learning Representations*, 2022.
- Kleinberg, J., Mullainathan, S., and Raghavan, M. Inherent trade-offs in the fair determination of risk scores. In Papadimitriou, C. H. (ed.), *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (Lipics)*, pp. 43:1–43:23, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-029-3.
- Kong, I., Kim, K., and Kim, Y. Fair Representation Learning for Continuous Sensitive Attributes Using Expectation of Integral Probability Metrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(5):3784–3795, May 2025. ISSN 1939-3539.
- Kwak, N. Nonlinear Projection Trick in Kernel Methods: An Alternative to the Kernel Trick. *IEEE Transactions on Neural Networks and Learning Systems*, 24(12):2113–2119, 2013.
- Lai, Y. and Guan, L. FairICP: Encouraging equalized odds via inverse conditional permutation. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 32228–32245. PMLR, July 2025.

- 495 Mary, J., Calauzènes, C., and Karoui, N. E. Fairness-Aware
 496 Learning for Continuous Attributes and Treatments. In
 497 *Proceedings of the 36th International Conference on Ma-*
 498 *chine Learning*, pp. 4382–4391. PMLR, May 2019.
- 499 Narasimhan, H., Cotter, A., Gupta, M., and Wang, S. L.
 500 Pairwise fairness for ranking and regression. In *33rd*
 501 *AAAI Conference on Artificial Intelligence*, June 2020.
- 503 Olfat, M. and Aswani, A. Spectral Algorithms for Comput-
 504 ing Fair Support Vector Machines. In Storkey, A. and
 505 Perez-Cruz, F. (eds.), *Proceedings of the Twenty-First*
 506 *International Conference on Artificial Intelligence and*
 507 *Statistics*, volume 84 of *Proceedings of Machine Learning*
 508 *Research*, pp. 1933–1942. PMLR, April 2018.
- 510 Oneto, L., Donini, M., and Pontil, M. General fair em-
 511 pirical risk minimization. In *2020 International Joint*
 512 *Conference on Neural Networks (IJCNN)*, pp. 1–8, 2020.
- 514 Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V.,
 515 Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P.,
 516 Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cour-
 517 napeau, D., Brucher, M., Perrot, M., and Duchesnay, E.
 518 Scikit-learn: Machine Learning in Python. *Journal of*
 519 *Machine Learning Research*, 12:2825–2830, 2011.
- 520 Pérez-Suay, A., Laparra, V., Mateo-García, G., Muñoz-
 521 Marí, J., Gómez-Chova, L., and Camps-Valls, G. Fair
 522 Kernel Learning. In Ceci, M., Hollmén, J., Todorovski,
 523 L., Vens, C., and Džeroski, S. (eds.), *Machine Learning*
 524 *and Knowledge Discovery in Databases*, pp. 339–355,
 525 Cham, 2017. Springer International Publishing. ISBN
 526 978-3-319-71249-9.
- 528 Perez-Suay, A., Gordaliza, P., Loubes, J.-M., Sejdinovic,
 529 D., and Camps-Valls, G. Fair Kernel Regression through
 530 Cross-Covariance Operators. *Transactions on Machine*
 531 *Learning Research*, 2023. ISSN 2835-8856.
- 533 Rasmussen, C. E. and Williams, C. K. I. *Gaussian Pro-*
 534 *cesses for Machine Learning*. Adaptive Computation
 535 and Machine Learning. MIT Press, Cambridge, Mass., 3.
 536 print edition, 2008. ISBN 978-0-262-18253-9.
- 537 Ravfogel, S., Elazar, Y., Gonen, H., Twiton, M., and Gold-
 538 berg, Y. Null It Out: Guarding Protected Attributes by
 539 Iterative Nullspace Projection. In Jurafsky, D., Chai,
 540 J., Schluter, N., and Tetreault, J. (eds.), *Proceedings of*
 541 *the 58th Annual Meeting of the Association for Compu-*
 542 *tational Linguistics*, pp. 7237–7256, Online, July 2020.
 543 Association for Computational Linguistics.
- 545 Redmond, M. and Baveja, A. A data-driven software tool for
 546 enabling cooperative information sharing among police
 547 departments. *European Journal of Operational Research*,
 548 141(3):660–678, 2002. ISSN 0377-2217.
- Scholkopf, B., Mika, S., Burges, C., Knirsch, P., Muller, K.-
 R., Ratsch, G., and Smola, A. Input space versus feature
 space in kernel-based methods. *IEEE Transactions on*
Neural Networks, 10(5):1000–1017, 1999.
- Schölkopf, B., Herbrich, R., and Smola, A. J. A Generalized
 Representer Theorem. In Helmbold, D. and Williamson,
 B. (eds.), *Computational Learning Theory*, pp. 416–426,
 Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
 ISBN 978-3-540-44581-4.
- Shawe-Taylor, J. *Kernel Methods for Pattern Analysis*. Cam-
 bridge University Press, Cambridge, 2004. ISBN 978-0-
 511-21060-0.
- Stypinska, J. AI ageism: A critical roadmap for studying
 age discrimination and exclusion in digitalized societies.
AI & SOCIETY, 38(2):665–677, April 2023. ISSN 1435-
 5655.
- Tan, Z., Yeom, S., Fredrikson, M., and Talwalkar, A. Learn-
 ing fair representations for kernel models. In Chiappa,
 S. and Calandra, R. (eds.), *Proceedings of the Twenty*
Third International Conference on Artificial Intelligence
and Statistics, volume 108 of *Proceedings of Machine*
Learning Research, pp. 155–166. PMLR, August 2020.
- Union, E. Charter of Fundamental Rights of the European
 Union, October 2012.
- Williams, C. and Seeger, M. Using the Nyström Method
 to Speed Up Kernel Machines. In Leen, T., Dietterich,
 T., and Tresp, V. (eds.), *Advances in Neural Information*
Processing Systems, volume 13. MIT Press, 2000.
- Xiong, H., Swamy, M., and Ahmad, M. Optimizing the
 kernel in the empirical feature space. *IEEE Transactions*
on Neural Networks, 16(2):460–474, 2005.
- Xu, W., Smart, M., Tilahun, N., Askari, S., Dennis, Z., Li,
 H., and Levinson, D. The racial composition of road
 users, traffic citations, and police stops. *Proceedings of*
the National Academy of Sciences of the United States of
America, 121(24):e2402547121, 2024.
- Zafar, M. B., Valera, I., Gomez-Rodriguez, M., and Gum-
 madi, K. P. Fairness constraints: A flexible approach for
 fair classification. *Journal of Machine Learning Research*,
 20(75):1–42, 2019.
- Zare, H., Ponce, A. N., Valek, R., Masoudi, N., Webster,
 D., Jr, R. J. T., Spencer, M., Crifasi, C., and Gaskin, D.
 Analyzing Fatal Police Shootings: The Roles of Social
 Vulnerability, Race, and Place in the U.S. *American*
Journal of Preventive Medicine, 68(1):126–136, 2025.
 ISSN 0749-3797.

A. Proofs

A.1. Proof for Lemma 3.1:

Proof. Following the proof in Ravfogel et al. (2020, Appendix A) we only need to show that we have $\mathbf{w}_1^\top \mathbf{w}_2 = 0$ for two consecutive iterations of our algorithm. After finding \mathbf{w}_1 with Equation (6) and projection in Equation (8) we obtain $\mathbf{G}_{(1)}$ such that $\mathbf{G}_{(1)} \mathbf{w}_1 = 0$ (*). Iterating we now find \mathbf{w}_2 using $\mathbf{G}_{(1)}$. By the representer theorem (Schölkopf et al., 2001) this can be written as a finite linear combination $\mathbf{w}_2 = \sum_{i=1}^n \xi_i \mathbf{g}_i$ where the \mathbf{g}_i correspond to the rows in $\mathbf{G}_{(1)}$. By (*) it holds: $\mathbf{g}_i^\top \mathbf{w}_1 = \mathbf{w}_1^\top \mathbf{g}_i = 0$ and therefore:

$$\mathbf{w}_1^\top \mathbf{w}_2 = \mathbf{w}_1^\top \sum_{i=1}^n \xi_i \mathbf{g}_i = \sum_{i=1}^n \xi_i \mathbf{w}_1^\top \mathbf{g}_i = 0.$$

The full result follows by the same argument and corollaries that are found in Ravfogel et al. (2020, A.1.1-A.1.3). \square

A.2. Proof for Theorem 3.2:

Proof. Let $\mathbf{P}_* := (\text{Id} - \mathbf{w}(\mathbf{w}^\top \mathbf{w})^{-1} \mathbf{w}^\top)$, then \mathbf{P}_* is a projection matrix so it holds:

$$\begin{aligned} \mathbf{K}_{(m)} &= \mathbf{G}_{(m-1)} \mathbf{G}_{(m-1)}^\top \\ &= \mathbf{G}_{(m-1)} \mathbf{P}_* (\mathbf{G}_{(m-1)} \mathbf{P}_*)^\top \\ &= \mathbf{G}_{(m-1)} \mathbf{P}_* \mathbf{P}_*^\top \mathbf{G}_{(m-1)}^\top \\ &= \mathbf{G}_{(m-1)} \mathbf{P}_* \mathbf{G}_{(m-1)}^\top \end{aligned}$$

where we used the properties of a projection matrix $\mathbf{P}_*^\top = \mathbf{P}_*$ and $\mathbf{P}_*^2 = \mathbf{P}_*$. Now simplifying \mathbf{P}_* we can rewrite (for easier notation we omit the index $m-1$ here):

$$\begin{aligned} \mathbf{w} \tau_{norm} \mathbf{w}^\top &= \mathbf{G}^\top (\mathbf{G} \mathbf{G}^\top + \tilde{\alpha} \text{Id})^{-1} \mathbf{p} \tau_{norm} (\mathbf{G}^\top (\mathbf{G} \mathbf{G}^\top + \tilde{\alpha} \text{Id})^{-1} \mathbf{p})^\top \\ &= \mathbf{G}^\top (\mathbf{G} \mathbf{G}^\top + \tilde{\alpha} \text{Id})^{-1} \mathbf{p} \tau_{norm} \mathbf{p}^\top (\mathbf{G} \mathbf{G}^\top + \tilde{\alpha} \text{Id})^{-1} \mathbf{G} \\ &= \mathbf{G}^\top (\mathbf{K} + \tilde{\alpha} \text{Id})^{-1} \mathbf{p} \tau_{norm} \mathbf{p}^\top (\mathbf{K} + \tilde{\alpha} \text{Id})^{-1} \mathbf{G} \\ &= \mathbf{G}^\top \mathbf{M} \mathbf{G} \end{aligned}$$

where we define $\mathbf{M} := (\mathbf{K} + \tilde{\alpha} \text{Id})^{-1} \mathbf{p} \tau_{norm} \mathbf{p}^\top (\mathbf{K} + \tilde{\alpha} \text{Id})^{-1}$ and denote the normalization scalar by τ_{norm} and write it as

$$\tau_{norm} = (\mathbf{w}^\top \mathbf{w})^{-1} = \mathbf{p}^\top (\mathbf{K} + \tilde{\alpha} \text{Id})^{-1} \mathbf{K} (\mathbf{K} + \tilde{\alpha} \text{Id})^{-1} \mathbf{p}.$$

Further we have (again omitting the index):

$$\mathbf{G} \mathbf{P}_* \mathbf{G}^\top = \mathbf{G} (\text{Id} - \mathbf{G}^\top \mathbf{M} \mathbf{G}) \mathbf{G}^\top = (\mathbf{G} - \mathbf{G} \mathbf{G}^\top \mathbf{M} \mathbf{G}) \mathbf{G}^\top = (\mathbf{G} \mathbf{G}^\top - \mathbf{G} \mathbf{G}^\top \mathbf{M} \mathbf{G} \mathbf{G}^\top) =: (*)$$

hence with $\mathbf{K}_{(m-1)} = \mathbf{G}_{(m-1)} \mathbf{G}_{(m-1)}^\top$ we have

$$(*) = \mathbf{K}_{(m-1)} - \mathbf{K}_{(m-1)} \mathbf{M} \mathbf{K}_{(m-1)} = \mathbf{K}_{(m-1)} (\text{Id} - \mathbf{M} \mathbf{K}_{(m-1)})$$

which was to be shown. \square

This setup can be naturally extended to the case of multiple protected attributes by instead letting $\mathbf{W} \in \mathbb{R}^{n \times l}$ aiming to predict $\mathbf{p} \in \mathbb{R}^{n \times l}$. Then instead of the scalar τ_{norm} for the normalization we have a matrix $\mathbb{R}^{l \times l} \ni \mathcal{T} = (\mathbf{W}^\top \mathbf{W})^{-1}$. Therefore Theorem 3.2 naturally transfers.

A.3. Proof for Corollary 3.3:

Proof. By Equation (10) and Theorem 3.2 $\mathbf{K}_{(m)}$ is equivalent to the linear kernel applied to $\mathbf{G}_{(m)}$. \square

Alg.	Dataset	Parameters
KRR-FKD (ours)	All	$\alpha = 0.25, \gamma = 0.05, \tilde{\alpha} = 0.1$
KRR-FKL	All	Equal to KRR-FKD, $\tilde{\alpha}$ not used.
SVR-FKD (ours)	Crimes ACSIInc ACSTra All	$\epsilon = 0.01, \gamma = 0.05, C = 0.75$ $\epsilon = 0.005, \gamma = 0.05, C = 0.5$ $\epsilon = 0.001, \gamma = 0.01, C = 0.125$ $\tilde{\alpha} = 0.05$
NN-HGR	All	Adam, $\rho = 0.01, \eta = 1e - 3$, SELU, Top.: $n \times 100 \times 80 \times 1$ Training for 500 epochs.

Table A.1. Fixed hyper-parameters for the considered approaches. RBF parameter γ , ridge penalty α , ridge penalty for fairness $\tilde{\alpha}$, epsilon bandwidth ϵ , learning rate η and weight decay ρ . n inputs. Adam optimizer and SELU activation.

Alg.	Dataset	Regularization
KRR-FKD (ours)	Crimes ACSIInc ACSTra	$m \in \{0, 5, 10, 18\}$ $m \in \{0, 15, 20, 30\}$ $m \in \{0, 15, 20, 30\}$
KRR-FKL	Crimes ACSIInc ACSTra	$\mu \in \{0, 0.05, 0.25, 0.6, 1\}$ $\mu \in \{0, 0.05, 0.25, 1\}$ $\mu \in \{0, 0.5, 1, 2\}$
SVR-FKD (ours)	Crimes ACSIInc ACSTra	$m \in \{0, 5, 30, 45, 60, 80\}$ $m \in \{0, 45, 60, 80, 100\}$ $m \in \{0, 45, 60, 100\}$
NN-HGR	Crimes ACSIInc ACSTra	$\lambda \in \{0, 0.025, 0.1, 0.2, 0.5\}$ $\lambda \in \{0, 0.00625, 0.0125, 0.05\}$ $\lambda \in \{0, 0.00625, 0.0125, 0.05\}$

Table A.2. Fairness regularization parameters that controls the Fairness-Accuracy-Tradeoff. Number of iterations m and penalty coefficients λ, μ .

B. Experimental Details

For the "NN-HGR" we use the implementation provided by the original paper (Mary et al., 2019). The Support Vector and Kernel Ridge Regression used for training our approach with the modified kernel utilize scikit-learn (Pedregosa et al., 2011). Table A.1 lists the fixed hyper-parameters that do not directly control the degree of fairness regularization. The different used fairness parameters to attain the results shown in Figure 1 are listed in Table A.2.

C. Nystroem Approximation

The nystroem approximation of a kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ works by sampling a number of columns $p \ll n$ from \mathbf{K} . The indices of the columns are also referred to as "landmarks". With these we get the approximation (Williams & Seeger, 2000)

$$\mathbf{K} \approx \mathbf{K}_{np} \mathbf{K}_{pp}^{-1} \mathbf{K}_{pn}$$

which is exact if $p = n$ components are used or if \mathbf{K} is of rank p (and p linearly independent columns were chosen). \mathbf{K}_{pp} corresponds the intersection of the p chosen rows and columns, \mathbf{K}_{np} to the p columns which are n dimensional each. The matrix inversion lemma is given by (Rasmussen & Williams, 2008, p. 201)

$$\begin{aligned} (\mathbf{Z} + \mathbf{U} \mathbf{W} \mathbf{V}^\top)^{-1} &= \\ \mathbf{Z}^{-1} - \mathbf{Z}^{-1} \mathbf{U} (\mathbf{W}^{-1} + \mathbf{V}^\top \mathbf{Z}^{-1} \mathbf{U})^{-1} \mathbf{V}^\top \mathbf{Z}^{-1} \end{aligned}$$

which with the following definitions $\mathbf{Z} := \alpha \mathbf{I}$, $\mathbf{U} := \mathbf{K}_{np}$, $\mathbf{V} := \mathbf{K}_{pn}$, $\mathbf{W} := \mathbf{K}_{pp}^{-1}$ yields:

$$\begin{aligned} (\alpha \text{Id} + \mathbf{K})^{-1} &\approx (\alpha \text{Id} + \mathbf{K}_{np} \mathbf{K}_{pp}^{-1} \mathbf{K}_{pn})^{-1} \\ &= (\mathbf{Z} + \mathbf{U} \mathbf{W} \mathbf{V}^\top)^{-1} \\ &= \frac{1}{\alpha} \text{Id} - \frac{1}{\alpha^2} \mathbf{K}_{np} (\mathbf{K}_{pp} + \frac{1}{\alpha} \mathbf{K}_{pn} \mathbf{K}_{np})^{-1} \mathbf{K}_{pn}. \end{aligned}$$

Evidently this only requires to invert a $p \times p$ matrix which is in $\mathcal{O}(p^3)$. Note that $\mathbf{K}_{pn} = \mathbf{K}_{np}^\top$. Of course this is an approximation so numerical issues might occur, to prevent some of them we first perform the approximation and then return the symmetrized result – see Algorithm 2.

Further, by choosing the optimal order of matrix multiplication we end up with an additional $\mathcal{O}(n^2(p + l))$ (l protected attributes) in matrix multiplications if the transformations $\mathbf{T}_{(i)}$ are not stored explicitly. This is detailed in the code.

Algorithm 2 Nystroem Approx. for Matrix Inversion

- 1: **Input:** Kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$, reg. param. α
- 2: Draw p column/landmark indices $c \leftarrow [c_1, \dots, c_p]$
- 3: Set $\mathbf{K}_{np} \in \mathbb{R}^{n \times p}$ with entries $\mathbf{K}_{np}[i, j] \leftarrow \mathbf{K}[i, c_j]$
- 4: Set $\mathbf{K}_{pp} \in \mathbb{R}^{p \times p}$ with entries $\mathbf{K}_{pp}[i, j] \leftarrow \mathbf{K}[c_i, c_j]$
- 5: $\mathbf{A} \leftarrow (\mathbf{K}_{pp} + \frac{1}{\alpha} \mathbf{K}_{pn} \mathbf{K}_{np})^{-1}$ $\{p \times p \text{ inverse}\}$
- 6: $\mathbf{B} \leftarrow \frac{1}{\alpha} \text{Id} - \frac{1}{\alpha^2} \mathbf{K}_{np} \mathbf{A} \mathbf{K}_{pn}$
- 7: $\mathbf{B} \leftarrow \frac{\mathbf{B} + \mathbf{B}^\top}{2}$ $\{\text{ensure symmetry for numerical stability}\}$
- 8: **Output:** \mathbf{B} the approx. of the inverse $(\mathbf{K} + \alpha \text{Id})^{-1}$

D. Additional Experiments

In our experiments we also bench-marked against the most recent method (Kong et al., 2025) for continuous fairness which we denote by “NN-FREM” but found it to be surpassed by the other baselines (“KRR-FKL”), see Figure 4. We used the same fixed network parameters (Table A.1) as for “NN-HGR” and using the following method-specific parameters: $\gamma = 0.05$, $\sigma = 1$, $\lambda \in \{0, 0.5, 1.5, 5, 10\}$ across all datasets.

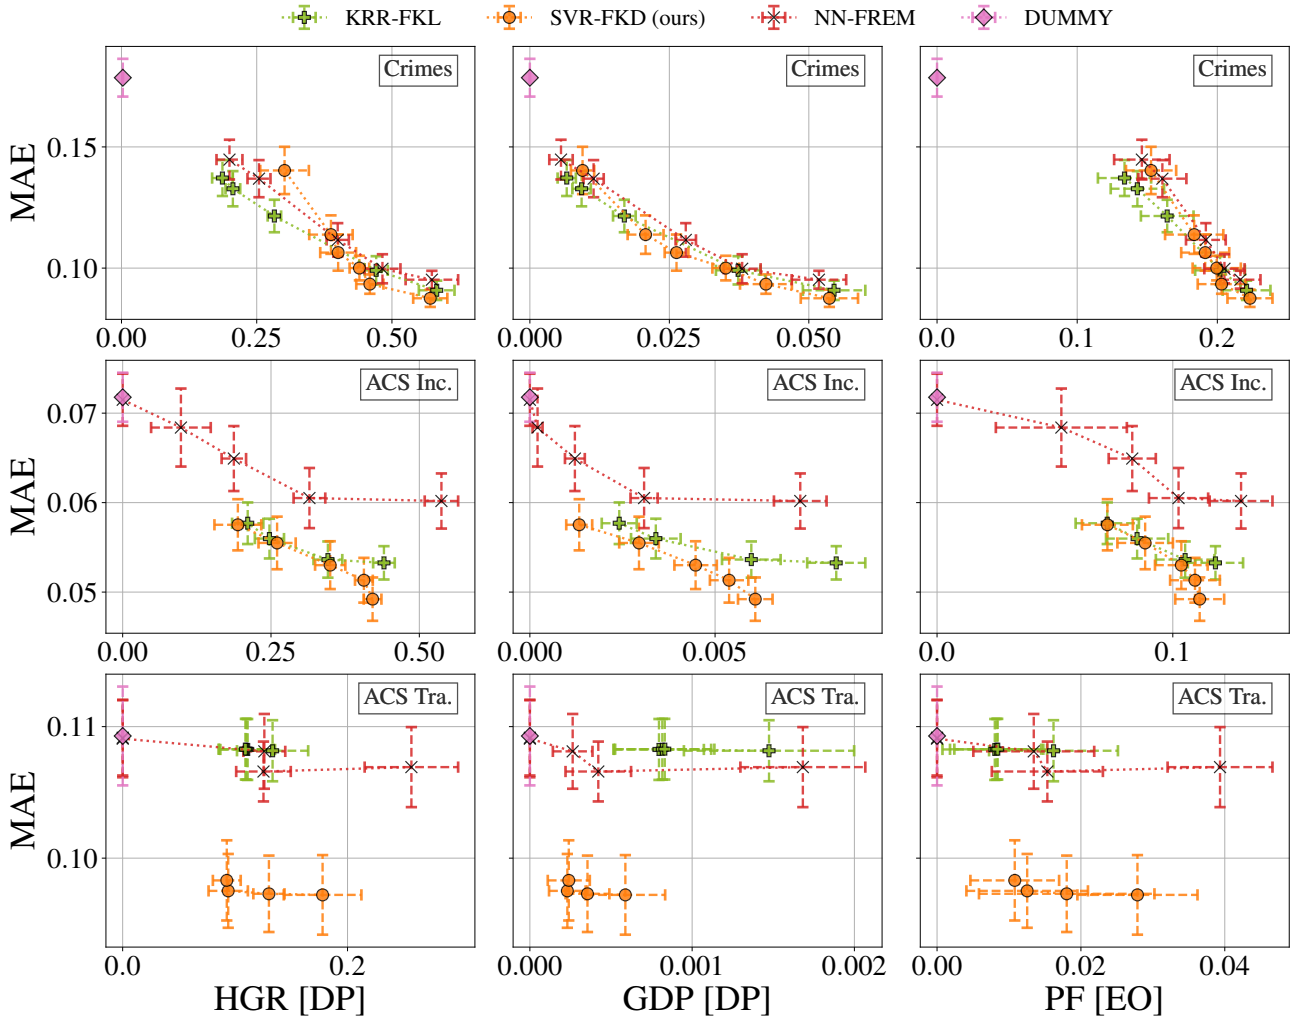


Figure 4. Same experimental setup as in Figure 1. Selected approaches compared to the most recent “NN-FREM” by Kong et al. (2025).