# Optimality Implies Kernel Sum Classifiers are Statistically Efficient 

Raphael A. Meyer ${ }^{1}$ Jean Honorio ${ }^{1}$

${ }^{1}$ Dept. of Computer Science, Purdue University

## Generalization Error Proofs

© Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

## Generalization Error Proofs

© Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathrm{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

© Prove the empirical error is close to the expected error

## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

© Prove the empirical error is close to the expected error

- Rademacher Complexity, PAC Bayes, etc.


## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

© Prove the empirical error is close to the expected error

- Rademacher Complexity, PAC Bayes, etc.
- Prove this for all $f \in \mathcal{F}$


## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

© Prove the empirical error is close to the expected error

- Rademacher Complexity, PAC Bayes, etc.
- Prove this for all $f \in \mathcal{F}$

$$
\underset{\mathbf{x}, y}{\mathbb{E}}[\ell(f(\mathrm{x}), y)] \leq \frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\mathbf{x}_{i}\right), y_{i}\right)+\varepsilon
$$

## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

© Prove the empirical error is close to the expected error

- Rademacher Complexity, PAC Bayes, etc.
- Prove this for all $f \in \mathcal{F}$

$$
\underset{\mathbf{x}, y}{\mathbb{E}}[\ell(f(\mathbf{x}), y)] \leq \frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\mathbf{x}_{i}\right), y_{i}\right)+\varepsilon
$$

## Generalization Error Proofs

(0) Assume we have $n$ i.i.d. labeled data samples

$$
\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)
$$

$$
\mathbf{x} \in \mathcal{X}, \quad y \in\{-1,1\}
$$

© Define a set of possible estimators to map inputs to labels

$$
\mathcal{F} \subseteq\{f: \mathcal{X} \mapsto\{-1,1\}\}
$$

© Prove the empirical error is close to the expected error

- Rademacher Complexity, PAC Bayes, etc.
- Prove this for all $f \in \mathcal{F}$
- This includes low-accuracy estimators $f$


## Optimization Theory

© Optimization is often used to find estimators in ML

## Optimization Theory

© Optimization is often used to find estimators in ML

- Regression, Kernel SVM, etc.


## Optimization Theory

© Optimization is often used to find estimators in ML - Regression, Kernel SVM, etc.
© These tools are used in practice

## Optimization Theory

© Optimization is often used to find estimators in ML - Regression, Kernel SVM, etc.
© These tools are used in practice
© Strong theoretical guarantees in polynomial time

## Optimization Theory

© Optimization is often used to find estimators in ML

- Regression, Kernel SVM, etc.
© These tools are used in practice
© Strong theoretical guarantees in polynomial time
- Karush-Kuhn-Tucker (KKT) Conditions


## Optimization Theory

© Optimization is often used to find estimators in ML

- Regression, Kernel SVM, etc.
© These tools are used in practice
© Strong theoretical guarantees in polynomial time
- Karush-Kuhn-Tucker (KKT) Conditions
© How can considering optimal estimators help us understand statistical efficiency?


## Multiple Kernel Learning \& Classification

Given:

## Multiple Kernel Learning \& Classification

Given:
© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, \mathrm{y}_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$

## Multiple Kernel Learning \& Classification

## Given:

© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$
© $m$ kernels $k_{1}, \ldots, k_{m}$

## Multiple Kernel Learning \& Classification

Given:
© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$
© $m$ kernels $k_{1}, \ldots, k_{m}$
Learn:

## Multiple Kernel Learning \& Classification

Given:
© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$
© $m$ kernels $k_{1}, \ldots, k_{m}$
Learn:
© Linear combination $\boldsymbol{\theta}$ of kernels

## Multiple Kernel Learning \& Classification

Given:
© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$
© $m$ kernels $k_{1}, \ldots, k_{m}$

## Learn:

(0) Linear combination $\boldsymbol{\theta}$ of kernels s.t. the resulting kernel

$$
k_{\Sigma}(\cdot, \cdot):=\sum_{t=1}^{m} \theta_{t} k_{t}(\cdot, \cdot)
$$

Classifies the dataset well

## Multiple Kernel Learning \& Classification

Given:
© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$
© $m$ kernels $k_{1}, \ldots, k_{m}$

## Learn:

(0) Linear combination $\boldsymbol{\theta}$ of kernels s.t. the resulting kernel

$$
k_{\Sigma}(\cdot, \cdot):=\sum_{t=1}^{m} \theta_{t} k_{t}(\cdot, \cdot)
$$

Classifies the dataset well
(c) Constraints on $\boldsymbol{\theta}$ vary among papers

## Multiple Kernel Learning \& Classification

Given:
© Dataset with $n$ samples $\left\{\left(\mathrm{x}_{1}, y_{1}\right), \ldots,\left(\mathrm{x}_{n}, y_{n}\right)\right\}$
© $m$ kernels $k_{1}, \ldots, k_{m}$

## Learn:

(0) Linear combination $\boldsymbol{\theta}$ of kernels s.t. the resulting kernel

$$
k_{\Sigma}(\cdot, \cdot):=\sum_{t=1}^{m} \theta_{t} k_{t}(\cdot, \cdot)
$$

Classifies the dataset well
© Constraints on $\boldsymbol{\theta}$ vary among papers

- $\boldsymbol{\theta}$ may be convex combination, $0 / 1$ vector, etc.
( Estimators are uniquely identified by $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ and $\boldsymbol{\theta}$

$$
f(\mathbf{x} ; \boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i} y_{i} k_{\Sigma}\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Prior Work: Generalization Error

(0) Estimators are uniquely identified by $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ and $\boldsymbol{\theta}$

$$
f(\mathbf{x} ; \boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i} y_{i} k_{\Sigma}\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Theorem [CMR10]

For all kernels, vectors $\boldsymbol{\alpha}$, and convex combinations $\boldsymbol{\theta}$ where

## Prior Work: Generalization Error

(0) Estimators are uniquely identified by $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ and $\boldsymbol{\theta}$

$$
f(\mathbf{x} ; \boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i} y_{i} k_{\Sigma}\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Theorem [CMR10]

For all kernels, vectors $\boldsymbol{\alpha}$, and convex combinations $\boldsymbol{\theta}$ where
( $k_{t}\left(\mathbf{x}_{i}, \mathbf{x}_{i}\right) \leq R^{2}$ for all $\mathbf{x}_{i}$ and $k_{t}$

## Prior Work: Generalization Error

(0) Estimators are uniquely identified by $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ and $\boldsymbol{\theta}$

$$
f(\mathbf{x} ; \boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i} y_{i} k_{\Sigma}\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Theorem [CMR10]

For all kernels, vectors $\boldsymbol{\alpha}$, and convex combinations $\boldsymbol{\theta}$ where
( $k_{t}\left(\mathbf{x}_{i}, \mathbf{x}_{i}\right) \leq R^{2}$ for all $\mathbf{x}_{i}$ and $k_{t}$
() $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha} \leq C^{2}$ for kernel matrix $\tilde{\boldsymbol{K}}_{\Sigma}=\sum_{t=1}^{m} \theta_{t} \tilde{\boldsymbol{K}}_{t}$

## Prior Work: Generalization Error

( Estimators are uniquely identified by $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ and $\boldsymbol{\theta}$

$$
f(\mathbf{x} ; \boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i} y_{i} k_{\Sigma}\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Theorem [CMR10]

For all kernels, vectors $\boldsymbol{\alpha}$, and convex combinations $\boldsymbol{\theta}$ where
( $k_{t}\left(\mathrm{x}_{i}, \mathbf{x}_{i}\right) \leq R^{2}$ for all $\mathrm{x}_{i}$ and $k_{t}$
() $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha} \leq C^{2}$ for kernel matrix $\tilde{\boldsymbol{K}}_{\Sigma}=\sum_{t=1}^{m} \theta_{t} \tilde{\boldsymbol{K}}_{t}$

We have

$$
\underset{\mathrm{x}, \mathrm{y}}{\mathbb{E}}[\ell(f(\mathrm{x} ; \alpha), y)] \leq \frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\mathrm{x}_{i} ; \alpha\right), y_{i}\right)+O\left(\frac{C R}{\sqrt{n}} \sqrt{\ln m}\right)
$$

## Prior Work: Generalization Error

( Estimators are uniquely identified by $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ and $\boldsymbol{\theta}$

$$
f(\mathbf{x} ; \boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i} y_{i} k_{\Sigma}\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Theorem [CMR10]

For all kernels, vectors $\boldsymbol{\alpha}$, and convex combinations $\boldsymbol{\theta}$ where
( $k_{t}\left(\mathbf{x}_{i}, \mathbf{x}_{i}\right) \leq R^{2}$ for all $\mathbf{x}_{i}$ and $k_{t}$
() $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha} \leq C^{2}$ for kernel matrix $\tilde{\boldsymbol{K}}_{\Sigma}=\sum_{t=1}^{m} \theta_{t} \tilde{\boldsymbol{K}}_{t}$

We have

$$
\underset{\mathrm{x}, \mathrm{y}}{\mathbb{E}}[\ell(f(\mathrm{x} ; \alpha), y)] \leq \frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\mathrm{x}_{i} ; \alpha\right), y_{i}\right)+O\left(\frac{C R}{\sqrt{n}} \sqrt{\ln m}\right)
$$

## How can we understand $\alpha^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ ?

How does $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ depend on
© The number of kernels?

## How can we understand $\alpha^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \alpha$ ?

How does $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ depend on
© The number of kernels?
() The conditioning of the kernels?

## How can we understand $\alpha^{\top} \tilde{K}_{\Sigma} \alpha$ ?

How does $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ depend on
() The number of kernels?
() The conditioning of the kernels?

These questions are ill-posed:

## How can we understand $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ ?

How does $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ depend on
() The number of kernels?
© The conditioning of the kernels?
These questions are ill-posed:
( $)$ All $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ define a feasible estimator

## How can we understand $\alpha^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \alpha$ ?

How does $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ depend on
() The number of kernels?
© The conditioning of the kernels?
These questions are ill-posed:
(0) All $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ define a feasible estimator
() So there always exists a feasible estimator with large $\boldsymbol{\alpha}^{\top} \tilde{K} \boldsymbol{\alpha}$

## How can we understand $\alpha^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \alpha$ ?

How does $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}$ depend on
() The number of kernels?
© The conditioning of the kernels?
These questions are ill-posed:
(0) All $\boldsymbol{\alpha} \in \mathbb{R}^{n}$ define a feasible estimator
(© So there always exists a feasible estimator with large $\boldsymbol{\alpha}^{\top} \tilde{K} \boldsymbol{\alpha}$
© But Support Vector Machines pick $\boldsymbol{\alpha}$ in practice

## How can we understand $\alpha^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \alpha$ ?

If $\boldsymbol{\alpha}_{\Sigma}$ solves the SVM problem with $\tilde{\boldsymbol{K}}_{\Sigma}$,
How does $\boldsymbol{\alpha}_{\Sigma}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}_{\Sigma}$ depend on
© The number of kernels?
© The conditioning of the kernels?
These questions are ill-posed:
( ) All $\alpha \in \mathbb{R}^{n}$ define a feasible estimator
( ) So there always exists a feasible estimator with large $\alpha^{\top} \tilde{K} \alpha$
© But Support Vector Machines pick $\alpha$ in practice

## Our Approach

Given:
$\begin{array}{llllll}\tilde{\boldsymbol{K}}_{1} & \tilde{\boldsymbol{K}}_{2} & \cdots & \tilde{\boldsymbol{K}}_{m} & \tilde{\boldsymbol{K}}_{\bar{\Sigma}}\end{array}$

## Our Approach



## Our Approach



## Our Approach

Given:

Optimize:

Assume:


Then $\alpha_{\Sigma}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}_{\Sigma} \leq 3 m^{-0.58} B^{2}$

## Our Approach

Given:

Optimize:

Assume:


Then $\boldsymbol{\alpha}_{\Sigma}^{\top} \tilde{\boldsymbol{K}}_{\Sigma} \boldsymbol{\alpha}_{\Sigma} \leq 3 m^{-0.58} B^{2}$

Then:
Estimator $y\left(\mathbf{x} ; \boldsymbol{\alpha}_{\Sigma}\right)$ generalizes well

$$
O\left(\frac{B R m^{0.208} \sqrt{\ln m}}{\sqrt{n}}\right)
$$

## Conclusion

© Leverage Optimization Theory to ask and answer well-posed questions about the statistics of practical estimators.
© Leverage Optimization Theory to ask and answer well-posed questions about the statistics of practical estimators.
© Consider $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}} \boldsymbol{\alpha}$ in Multiple Kernel Learning as a specific case
© Leverage Optimization Theory to ask and answer well-posed questions about the statistics of practical estimators.
© Consider $\boldsymbol{\alpha}^{\top} \tilde{\boldsymbol{K}} \boldsymbol{\alpha}$ in Multiple Kernel Learning as a specific case
© Several possible applications of this idea beyond kernels

## THANK

Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Generalization Bounds for Learning Kernels.
In Proceedings of the 27th International Conference on International Conference on Machine Learning, pages 247-254.
Omnipress, 2010.
Raphael Meyer and Jean Honorio.
Optimality Implies Kernel Sum Classifiers are Statistically Efficient.

In International Conference on Machine Learning, pages 4566-4574, 2019.

