# Asymptotics of Gaussian Regularized Least-Squares 

Ross A. Lippert<br>M.I.T., Department of Mathematics<br>77 Massachusetts Avenue<br>Cambridge, MA 02139-4307<br>lippert@math.mit.edu

Ryan M. Rifkin<br>Honda Research Institute USA, Inc. 145 Tremont Street<br>Boston, MA 02111<br>rrifkin@honda-ri.com


#### Abstract

We consider regularized least-squares (RLS) with a Gaussian kernel. We prove that if we let the Gaussian bandwidth $\sigma \rightarrow \infty$ while letting the regularization parameter $\lambda \rightarrow 0$, the RLS solution tends to a polynomial whose order is controlled by the rielative rates of decay of $\frac{1}{\sigma^{2}}$ and $\lambda$ : if $\lambda=\sigma^{-(2 k+1)}$, then, as $\sigma \rightarrow \infty$, the RLS solution tends to the $k$ th order polynomial with minimal empirical error. We illustrate the result with an example.


## 1 Introduction

Given a data set $\left(x_{1}, y_{1}\right),\left(x_{2}, y_{2}\right), \ldots,\left(x_{n}, y_{n}\right)$, the inductive learning task is to build a function $f(x)$ that, given a new $x$ point, can predict the associated $y$ value. We study the Regularized Least-Squares (RLS) algorithm for finding $f$, a common and popular algorithm [2,5] that can be used for either regression or classification:

$$
\min _{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n}\left(f\left(x_{i}\right)-y_{i}\right)^{2}+\lambda\|f\|_{K}^{2}
$$

Here, $\mathcal{H}$ is a Reproducing Kernel Hilbert Space (RKHS) [1] with associated kernel function $K,\|f\|_{K}^{2}$ is the squared norm in the RKHS, and $\lambda$ is a regularization constant controlling the tradeoff between fitting the training set accurately and forcing smoothness of $f$.
The Representer Theorem [7] proves that the RLS solution will have the form $f(x)=$ $\sum_{i=1}^{n} c_{i} K\left(x_{i}, x\right)$, and it is easy to show [5] that we can find the coefficients $c$ by solving the linear system

$$
\begin{equation*}
(K+\lambda n I) c=y \tag{1}
\end{equation*}
$$

where $K$ is the $n$ by $n$ matrix satisfying $K_{i j}=K\left(x_{i}, x_{j}\right)$. We focus on the Gaussian kernel $K\left(x_{i}, x_{j}\right)=\exp \left(-\left\|x_{i}-x_{j}\right\|^{2} / 2 \sigma^{2}\right)$.
Our work was originally motivated by the empirical observation that on a range of benchmark classification tasks, we achieved surprisingly accurate classification using a Gaussian kernel with a very large $\sigma$ and a very small $\lambda$ (Figure 1; additional examples in [6]). This prompted us to study the large- $\sigma$ asymptotics of RLS. As $\sigma \rightarrow \infty, K\left(x_{i}, x_{j}\right) \rightarrow 1$ for arbitrary $x_{i}$ and $x_{j}$. Consider a single test point $x_{0}$. RLS will first find $c$ using Equation 1,

## RLSC Results for GALAXY Dataset



Fig. 1. RLS classification accuracy results for the UCI Galaxy dataset over a range of $\sigma$ (along the $x$-axis) and $\lambda$ (different lines) values. The vertical labelled lines show $m$, the smallest entry in the kernel matrix for a given $\sigma$. We see that when $\lambda=1 e-11$, we can classify quite accurately when the smallest entry of the kernel matrix is .99999 .
then compute $f\left(x_{0}\right)=c^{t} k$ where $k$ is the kernel vector, $k_{i}=K\left(x_{i}, x_{0}\right)$. Combining the training and testing steps, we see that $f\left(x_{0}\right)=y^{t}(K+\lambda n I)^{-1} k$.
Both $K$ and $k$ are close to 1 for large $\sigma$, i.e. $K_{i j}=1+\epsilon_{i j}$ and $k_{i}=1+\epsilon_{i}$. If we directly compute $c=(K+\lambda n I)^{-1} y$, we will tend to wash out the effects of the $\epsilon_{i j}$ term as $\sigma$ becomes large. If, instead, we compute $f\left(x_{0}\right)$ by associating to the right, first computing point affinities $(K+\lambda n I)^{-1} k$, then the $\epsilon_{i j}$ and $\epsilon_{j}$ interact meaningfully; this interaction is crucial to our analysis.
Our approach is to Taylor expand the kernel elements (and thus $K$ and $k$ ) in $1 / \sigma$, noting that as $\sigma \rightarrow \infty$, consecutive terms in the expansion differ enormously. In computing ( $K+$ $\lambda n I)^{-1} k$, these scalings cancel each other out, and result in finite point affinities even as $\sigma \rightarrow \infty$. The asymptotic affinity formula can then be "transposed" to create an alternate expression for $f\left(x_{0}\right)$. Our main result is that if we set $\sigma^{2}=s^{2}$ and $\lambda=s^{-(2 k+1)}$, then, as $s \rightarrow \infty$, the RLS solution tends to the $k$ th order polynomial with minimal empirical error.
The main theorem is proved in full. Due to space restrictions, the proofs of supporting lemmas and corollaries are omitted; an expanded version containing all proofs is available [4].

## 2 Notation and definitions

Definition 1. Let $x_{i}$ be a set of $n+1$ points $(0 \leq i \leq n)$ in a d dimensional space. The scalar $x_{i a}$ denotes the value of the $a^{\text {th }}$ vector component of the $i^{\text {th }}$ point.

The $n \times d$ matrix, $X$ is given by $X_{i a}=x_{i a}$.
We think of $X$ as the matrix of training data $x_{1}, \ldots, x_{n}$ and $x_{0}$ as an $1 \times d$ matrix consisting of the test point.

Let $1_{m}, 1_{l m}$ denote the $m$ dimensional vector and $l \times m$ matrix with components all 1 , similarly for $0_{m}, 0_{l m}$. We will dispense with such subscripts when the dimensions are clear from context.

Definition 2 (Hadamard products and powers). For two $l \times$ m matrices, $N, M, N \odot M$ denotes the $l \times m$ matrix given by $(N \odot M)_{i j}=N_{i j} M_{i j}$. Analogously, we set $\left(N^{\odot c}\right)_{i j}=$ $N_{i j}^{c}$.

Definition 3 (polynomials in the data). Let $I \in \mathbb{Z}_{\geq 0}^{d}$ (non-negative multi-indices) and $Y$ be a $k \times d$ matrix. $Y^{I}$ is the $k$ dimensional vector given by $\left(Y^{I}\right)_{i}=\prod_{a=1}^{d} Y_{i a}^{I_{a}}$. If $h: \mathbb{R}^{d} \rightarrow \mathbb{R}$ then $h(Y)$ is the $k$ dimensional vector given by $(h(Y))_{i}=h\left(Y_{i 1}, \ldots, Y_{i d}\right)$.
The $d$ canonical vectors, $e_{a} \in \mathbb{Z}_{\geq 0}^{d}$, are given by $\left(e_{a}\right)_{b}=\delta_{a b}$.
Any scalar function, $f: \mathbb{R} \rightarrow \mathbb{R}$, applied to any matrix or vector, $A$, will be assumed to denote the elementwise application of $f$. We will treat $y \rightarrow e^{y}$ as a scalar function (we have no need of matrix exponentials in this work, so the notation is unambiguous).
We can re-express the kernel matrix and kernel vector in this notation:

$$
\begin{align*}
K & \left.=e^{\frac{1}{2 \sigma^{2}} \sum_{a=1}^{d} 2 X^{e_{a}}\left(X^{e} a\right.}\right)^{t}-X^{2 e_{a}} 1_{n}^{t}-1_{n}\left(X^{2 e_{a}}\right)^{t}  \tag{2}\\
& =\operatorname{diag}\left(e^{-\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right) e^{\frac{1}{\sigma^{2}} X X^{t}} \operatorname{diag}\left(e^{-\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right)  \tag{3}\\
k & =e^{\frac{1}{2 \sigma^{2}} \sum_{a=1}^{d} 2 X^{e_{a}} x_{0}^{e_{a}}-X^{2 e_{a}} 1_{1}-1_{n} x_{0}^{2 e_{a}}}  \tag{4}\\
& =\operatorname{diag}\left(e^{-\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right) e^{\frac{1}{\sigma^{2}} X x_{0}^{t}} e^{-\frac{1}{2 \sigma^{2}}\left\|x_{0}\right\|^{2}} . \tag{5}
\end{align*}
$$

## 3 Orthogonal polynomial bases

Let $V_{c}=\operatorname{span}\left\{X^{I}:|I|=c\right\}$ and $V_{\leq c}=\bigcup_{a=0}^{c} V_{c}$ which can be thought of as the set of all $d$ variable polynomials of degree $c$, evaluated on the training data. Since the data are finite, there exists $b$ such that $V_{\leq c}=V_{\leq b}$ for all $c \geq b$. Generically, $b$ is the smallest $c$ such that $\binom{c+d}{d} \geq n$.

Let $Q$ be an orthonormal matrix in $\mathbb{R}^{n \times n}$ whose columns progressively span the $V_{\leq c}$ spaces, i.e. $Q=\left(\begin{array}{llll}B_{0} & B_{1} & \cdots & B_{b}\end{array}\right)$ where $Q^{t} Q=I$ and colspan $\left\{\left(\begin{array}{lll}B_{0} & \cdots & B_{c}\end{array}\right)\right\}=$ $V_{\leq c}$. We might imagine building such a $Q$ via the Gramm-Schmidt process on the vectors $\overline{X^{0}}, X^{e_{1}}, \ldots, X^{e_{d}}, \ldots X^{I}, \ldots$ taken in order of non-decreasing $|I|$.
Letting $C_{I}=\binom{|I|}{I_{1} \ldots I_{d}}$ be multinomial coefficients, the following relations between $Q, X$, and $x_{0}$ are easily proved.

$$
\begin{aligned}
\left(X x_{0}^{t}\right)^{\odot c} & =\sum_{|I|=c} C_{I} X^{I}\left(x_{0}^{I}\right)^{t} \quad \text { hence } \quad\left(X x_{0}^{t}\right)^{\odot c} \in V_{c} \\
\left(X X^{t}\right)^{\odot c} & =\sum_{|I|=c} C_{I} X^{I}\left(X^{I}\right)^{t} \quad \text { hence } \quad \text { colspan }\left\{\left(X X^{t}\right)^{\odot c}\right\}=V_{c}
\end{aligned}
$$

and thus, $B_{i}^{t}\left(X x_{0}^{t}\right)^{\odot c}=0$ if $i>c, B_{i}^{t}\left(X X^{t}\right)^{\odot c} B_{j}=0$ if $i>c$ or $j>c$, and $B_{c}^{t}\left(X X^{t}\right)^{\odot}{ }^{c} B_{c}$ is non-singular.
Finally, we note that $\operatorname{argmin}_{v \in V_{\leq c}}\{\|y-v\|\}=\sum_{a \leq c} B_{a}\left(B_{a}^{t} y\right)$.

## 4 Taking the $\sigma \rightarrow \infty$ limit

We will begin with a few simple lemmas about the limiting solutions of linear systems. At the end of this section we will arrive at the limiting form of suitably modified RLSC equations.

Lemma 1. Let $i_{1}<\cdots<i_{q}$ be positive integers. Let $A(s), y(s)$ be a block matrix and block vector given by

$$
A(s)=\left(\begin{array}{cccc}
A_{00}(s) & s^{i_{1}} A_{01}(s) & \cdots & s^{i_{q}} A_{0 q}(s) \\
s^{i_{1}} A_{10}(s) & s^{i_{1}} A_{11}(s) & \cdots & s^{i_{q}} A_{1 q}(s) \\
\cdots & \cdots & \cdots & \cdots \\
s^{i_{q}} A_{q 0}(s) & s^{i_{q}} A_{q 1}(s) & \cdots & s^{i_{q}} A_{q q}(s)
\end{array}\right), \quad y(s)=\left(\begin{array}{c}
b_{0}(s) \\
s^{i_{1}} b_{1}(s) \\
\cdots \\
s^{i_{q}} b_{q}(s)
\end{array}\right)
$$

where $A_{i j}(s)$ and $b_{i}(s)$ are continuous matrix-valued and vector-valued functions of $s$ with $A_{i i}(0)$ non-singular for all $i$.

$$
\lim _{s \rightarrow 0} A^{-1}(s) y(s)=\left(\begin{array}{cccc}
A_{00}(0) & 0 & \cdots & 0 \\
A_{10}(0) & A_{11}(0) & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots \\
A_{q 0}(0) & A_{q 1}(0) & \cdots & A_{q q}(0)
\end{array}\right)^{-1}\left(\begin{array}{c}
b_{0}(0) \\
b_{1}(0) \\
\cdots \\
b_{q}(0)
\end{array}\right)
$$

We are now ready to state and prove the main result of this section, characterizing the limiting large- $\sigma$ solution of Gaussian RLS.

Theorem 1. Let $q$ be an integer satisfying $q<b$, and let $p=2 q+1$. Let $\lambda=C \sigma^{-p}$ for some constant $C$. Define $A_{i j}^{(c)}=\frac{1}{c!} B_{i}^{t}\left(X X^{t}\right)^{\odot c} B_{j}$, and $b_{i}^{(c)}=\frac{1}{c!} B_{i}^{t}\left(X x_{0}^{t}\right)^{\odot c}$.

$$
\lim _{\sigma \rightarrow \infty}\left(K+n C \sigma^{-p} I\right)^{-1} k=v
$$

where

$$
\begin{align*}
v & =\left(\begin{array}{lll}
B_{0} & \cdots & B_{q}
\end{array}\right) w  \tag{6}\\
\left(\begin{array}{c}
b_{0}^{(0)} \\
b_{1}^{(1)} \\
\cdots \\
b_{q}^{(q)}
\end{array}\right) & =\left(\begin{array}{cccc}
A_{00}^{(0)} & 0 & \cdots & 0 \\
A_{10}^{(1)} & A_{11}^{(1)} & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots \\
A_{q 0}^{(q)} & A_{q 1}^{(q)} & \cdots & A_{q q}^{(q)}
\end{array}\right) w \tag{7}
\end{align*}
$$

We first manipulate the equation $(K+n \lambda I) y=k$ according to the factorizations in (3) and (5).

$$
\begin{aligned}
K & =\operatorname{diag}\left(e^{-\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right) e^{\frac{1}{\sigma^{2}} X X^{t}} \operatorname{diag}\left(e^{-\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right)=N P N \\
k & =\operatorname{diag}\left(e^{-\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right) e^{\frac{1}{\sigma^{2}} X x_{0}^{t}} e^{-\frac{1}{2 \sigma^{2}}\left\|x_{0}\right\|^{2}}=N w \alpha
\end{aligned}
$$

Noting that $\lim _{\sigma \rightarrow \infty} e^{-\frac{1}{2 \sigma^{2}}\left\|x_{0}\right\|^{2}} \operatorname{diag}\left(e^{\frac{1}{2 \sigma^{2}}\|X\|^{2}}\right)=\lim _{\sigma \rightarrow \infty} \alpha N^{-1}=I$, we have

$$
\begin{aligned}
v & \equiv \lim _{\sigma \rightarrow \infty}\left(K+n C \sigma^{-p} I\right)^{-1} k \\
& =\lim _{\sigma \rightarrow \infty}(N P N+\beta I)^{-1} N w \alpha \\
& =\lim _{\sigma \rightarrow \infty} \alpha N^{-1}\left(P+\beta N^{-2}\right)^{-1} w \\
& =\lim _{\sigma \rightarrow \infty}\left(e^{\frac{1}{\sigma^{2}} X X^{t}}+n C \sigma^{-p} \operatorname{diag}\left(e^{\frac{1}{\sigma^{2}}\|X\|^{2}}\right)\right)^{-1} e^{\frac{1}{\sigma^{2}} X x_{0}^{t}}
\end{aligned}
$$

Changing bases with $Q$,

$$
Q^{t} v=\lim _{\sigma \rightarrow \infty}\left(Q^{t} e^{\frac{1}{\sigma^{2}} X X^{t}} Q+n C \sigma^{-p} Q^{t} \operatorname{diag}\left(e^{\frac{1}{\sigma^{2}}\|X\|^{2}}\right) Q\right)^{-1} Q^{t} e^{\frac{1}{\sigma^{2}} X x_{0}^{t}}
$$

Expanding via Taylor series and writing in block form (in the $b \times b$ block structure of $Q$ ),

$$
\begin{gathered}
Q^{t} e^{\frac{1}{\sigma^{2}} X X^{t}} Q=Q^{t}\left(X X^{t}\right)^{\odot 0} Q+\frac{1}{1!\sigma^{2}} Q^{t}\left(X X^{t}\right)^{\odot 1} Q+\frac{1}{2!\sigma^{4}} Q^{t}\left(X X^{t}\right)^{\odot} Q+\cdots \\
=\left(\begin{array}{cccc}
A_{00}^{(0)} & 0 & \cdots & 0 \\
0 & 0 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots \\
0 & 0 & \cdots & 0
\end{array}\right)+\frac{1}{\sigma^{2}}\left(\begin{array}{ccc}
A_{00}^{(1)} & A_{01}^{(1)} & \cdots \\
A_{10}^{(1)} & A_{11}^{(1)} & \cdots \\
\cdots & \cdots & 0 \\
0 & 0 & \cdots \\
\cdots
\end{array}\right)+\cdots \\
Q^{t} e^{\frac{1}{\sigma^{2}} X x_{0}^{t}}=Q^{t}\left(X x_{0}^{t}\right)^{\odot 0}+\frac{1}{\sigma^{2}} Q^{t}\left(X x_{0}^{t}\right)^{\odot 1}+\frac{1}{\sigma^{4}} Q^{t}\left(X x_{0}^{t}\right)^{\odot 2}+\cdots \\
\\
=\left(\begin{array}{c}
b_{0}^{(0)} \\
0 \\
\cdots \\
0
\end{array}\right)+\frac{1}{\sigma^{2}}\left(\begin{array}{c}
b_{0}^{(1)} \\
b_{1}^{(1)} \\
\cdots \\
0
\end{array}\right)+\cdots \\
n C \sigma^{-p} Q^{t} \operatorname{diag}\left(e^{\frac{1}{\sigma^{2}}\|X\|^{2}}\right) Q=n C \sigma^{-p} I+\cdots .
\end{gathered}
$$

Since the $A_{c c}^{(c)}$ are non-singular, Lemma 3 applies, giving our result.

## 5 The classification function

When performing RLS, the actual prediction of the limiting classifier is given via

$$
f_{\infty}\left(x_{0}\right) \equiv \lim _{\sigma \rightarrow \infty} y^{t}\left(K+n C \sigma^{-p} I\right)^{-1} k
$$

Theorem 1 determines $v=\lim _{\sigma \rightarrow \infty}\left(K+n C \sigma^{-p} I\right)^{-1} k$,showing that $f_{\infty}\left(x_{0}\right)$ is a polynomial in the training data $X$. In this section, we show that $f_{\infty}\left(x_{0}\right)$ is, in fact, a polynomial in the test point $x_{0}$. We continue to work with the orthonormal vectors $B_{i}$ as well as the auxilliary quantities $A_{i j}^{(c)}$ and $b_{i}^{(c)}$ from Theorem 1.
Theorem 1 shows that $v \in V_{\leq q}$ : the point affinity function is a polynomial of degree $q$ in the training data, determined by (7).

$$
\begin{array}{rll}
\sum_{i, j \leq c} c!B_{i} A_{i j}^{(c)} B_{j}^{t}=\left(X X^{t}\right)^{\odot c} & \text { hence } \quad \sum_{j \leq c} c!B_{c} A_{c j}^{(c)} B_{j}^{t}=B_{c} B_{c}^{t}\left(X X^{t}\right)^{\odot c} \\
\sum_{i \leq c} c!B_{i} b_{i}^{(c)}=\left(X x_{0}^{t}\right)^{\odot c} & \text { hence } & c!B_{c} b_{i}^{(c)}=B_{c} B_{c}^{t}\left(X x_{0}^{t}\right)^{\odot c}
\end{array}
$$

we can restate Equation 7 in an equivalent form:

$$
\begin{align*}
&\left(\begin{array}{c}
B_{0}^{t} \\
\cdots \\
B_{q}^{t}
\end{array}\right)^{t}\left(\left(\begin{array}{c}
0!b_{0}^{(0)} \\
1!b_{1}^{(1)} \\
\cdots \\
q!b_{q}^{(q)}
\end{array}\right)-\left(\begin{array}{cccc}
0!A_{00}^{(0)} & 0 & \cdots & 0 \\
1!A_{10}^{(1)} & 1!A_{11}^{(1)} & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots \\
q!A_{q 0}^{(q)} & q!A_{q 1}^{(q)} & \cdots & q!A_{q q}^{(q)}
\end{array}\right)\left(\begin{array}{c}
B_{0}^{t} \\
\cdots \\
B_{q}^{t}
\end{array}\right) v\right)=0  \tag{8}\\
& \sum_{c \leq q} c!B_{c} b_{c}^{(c)}-\sum_{c \leq q} \sum_{j \leq c} c!B_{c} A_{c j}^{(c)} B_{j}^{t} v=0  \tag{9}\\
& \sum_{c \leq q} B_{c} B_{c}^{t}\left(\left(X x_{0}^{t}\right)^{\odot c}-\left(X X^{t}\right)^{\odot c} v\right)=0 \tag{10}
\end{align*}
$$

Up to this point, our results hold for arbitrary training data $X$. To proceed, we require a mild condition on our training set.

Definition 4. $X$ is called generic if $X^{I_{1}}, \ldots, X^{I_{n}}$ are linearly independent for any distinct multi-indices $\left\{I_{i}\right\}$.

Lemma 2. For generic $X$, the solution to Equation 7 (or equivalently, Equation 10) is determined by the conditions $\forall I:|I| \leq q,\left(X^{I}\right)^{t} v=x_{0}^{I}$, where $v \in V_{\leq q}$.

Theorem 2. For generic data, let $v$ be the solution to Equation 10. For any $y \in \mathbb{R}^{n}$, $f\left(x_{0}\right)=y^{t} v=h\left(x_{0}\right)$, where $h(x)=\sum_{|I| \leq q} a_{I} x^{I}$ is a multivariate polynomial of degree $q$ minimizing $\|y-h(X)\|$.

We see that as $\sigma \rightarrow \infty$, the RLS solution tends to the minimum empirical error $k$ th order polynomial.

## 6 Experimental Verification

In this section, we present a simple experiment that illustrates our results. We consider a fith-degree polynomial function. Figure 2 plots $f$, along with a 150 point dataset drawn by choosing $x_{i}$ uniformly in $[0,1]$, and choosing $y=f(x)+\epsilon_{i}$, where $\epsilon_{i}$ is a Gaussian random variable with mean 0 and standard deviation .05 . Figure 2 also shows (in red) the best polynomial approximations to the data (not to the ideal $f$ ) of various orders. (We omit third order because it is nearly indistinguishable from second order.)
According to Theorem 1, if we parametrize our system by a variable $s$, and solve a Gaussian regularized least-squares problem with $\sigma^{2}=s^{2}$ and $\lambda=C s^{-(2 k+1)}$ for some integer $k$, then, as $s \rightarrow \infty$, we expect the solution to the system to tend to the $k$ th-order databased polynomial approximation to $f$. Asymptotically, the value of the constant $C$ does not matter, so we (arbitrarily) set it to be 1 . Figure 3 demonstrates this result.
We note that these experiments frequently require setting $\lambda$ much smaller than machine$\epsilon$. As a consequence, we need more precision than IEEE double-precision floating-point, and our results cannot be obtained via many standard tools (e.g., MATLAB(TM)) We performed our experiments using CLISP, an implementation of Common Lisp that includes arithmetic operations on arbitrary-precision floating point numbers.

## 7 Discussion

Our result provides insight into the asymptotic behavior of RLS, and (partially) explains Figure 1: in conjunction with additional experiments not reported here, we believe that


Fig. 2. $f(x)=.5(1-x)+150 x(x-.25)(x-.3)(x-.75)(x-.95)$, a random dataset drawn from $f(x)$ with added Gaussian noise, and data-based polynomial approximations to $f$.
we are recovering second-order polynomial behavior, with the drop-off in performance at various $\lambda$ 's occurring at the transition to third-order behavior, which cannot be accurately recovered in IEEE double-precision floating-point. Although we used the specific details of RLS in deriving our solution, we expect that in practice, a similar result would hold for Support Vector Machines, and perhaps for Tikhonov regularization with convex loss more generally.
An interesting implication of our theorem is that for very large $\sigma$, we can obtain various order polynomial classifications by sweeping $\lambda$. In [6], we present an algorithm for solving for a wide range of $\lambda$ for essentially the same cost as using a single $\lambda$. This algorithm is not currently practical for large $\sigma$, due to the need for extended-precision floating point.

Our work also has implications for approximations to the Gaussian kernel. Yang et al. use the Fast Gauss Transform (FGT) to speed up matrix-vector multiplications when performing RLS [8]. In [6], we studied this work; we found that while Yang et al. used moderate-tosmall values of $\sigma$ (and did not tune $\lambda$ ), the FGT sacrificed substantial accuracy compared to the best achievable results on their datasets. We showed empirically that the FGT becomes much more accurate at larger values of $\sigma$; however, at large- $\sigma$, it seems likely we are merely recovering low-order polynomial behavior. We suggest that approximations to the Gaussian kernel must be checked carefully, to show that they produce sufficiently good results are moderate values of $\sigma$; this is a topic for future work.

## References

1. Aronszajn. Theory of reproducing kernels. Transactions of the American Mathematical Society, 68:337-404, 1950.
2. Evgeniou, Pontil, and Poggio. Regularization networks and support vector machines. Advances In Computational Mathematics, 13(1):1-50, 2000.


Fig. 3. As $s \rightarrow \infty, \sigma^{2}=s^{2}$ and $\lambda=s^{-(2 k+1)}$, the solution to Gaussian RLS approaches the $k$ th order polynomial solution.
3. Keerthi and Lin. Asymptotic behaviors of support vector machines with gaussian kernel. Neural Computation, 15(7):1667-1689, 2003.
4. Ross Lippert and Ryan Rifkin. Asymptotics of gaussian regularized least-squares. Technical Report MIT-CSAIL-TR-2005-067, MIT Computer Science and Artificial Intelligence Laboratory, 2005.
5. Rifkin. Everything Old Is New Again: A Fresh Look at Historical Approaches to Machine Learning. PhD thesis, Massachusetts Institute of Technology, 2002.
6. Rifkin and Lippert. Practical regularized least-squares: $\lambda$-selection and fast leave-one-outcomputation. In preparation, 2005.
7. Wahba. Spline Models for Observational Data, volume 59 of CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial \& Applied Mathematics, 1990.
8. Yang, Duraiswami, and Davis. Efficient kernel machines using the improved fast Gauss transform. In Advances in Neural Information Processing Systems, volume 16, 2004.

