Deep Neural Networks and Kernels
10716, Spring 2020
Pradeep Ravikumar

1 Recap: RKHS Kernel Regression

As we have seen, kernel methods (or kernel machines as they are sometimes called) learn functions in a Reproducing Kernel Hilbert Space (RKHS). Recall that an RKHS is specified by a Mercer kernel $K : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$, and loosely, consists of functions of the form $f(x) = \sum_i \alpha_i K(x_i, x)$, where $\{x_i\} \subseteq \mathcal{X}$. More formally, consider the function space:

$$\mathcal{H}_0 = \{f \ : \ f(x) = \sum_{j=1}^{m} \alpha_j K(x_j, x), \ \alpha_j \in \mathbb{R}, x_j \in \mathcal{X}, m \in \mathbb{N}\}.$$

Given two such functions $f = \sum_i \alpha_i K(x_i, x)$ and $g(x) = \sum_j \beta_j K(y_j, x)$, we can define the inner product:

$$\langle f, g \rangle_K = \sum_i \sum_j \alpha_i \beta_j K(x_i, y_j),$$

which then defines a norm $\|f\|_K = \sqrt{\langle f, f \rangle_K}$. The RKHS $\mathcal{H}_K$ corresponding to the kernel K is then the completion of $\mathcal{H}_0$ with respect to the inner product $\| \cdot \|_K$.

(Ridge) RKHS regression, also simply called kernel (ridge) regression in the ML literature (not to be confused with Nadaraya Watson smoothing kernel based non-parametric regression), then consists of solving:

$$\inf_{f \in \mathcal{H}_K} \left\{ \widehat{R}_n(f) + \lambda \|f\|_K^2 \right\},$$

where $\widehat{R}_n(f) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(x_i), y_i)$ is some empirical risk. A typical loss function is the squared loss $\mathcal{L}(f(x_i), y_i) = \frac{1}{2}(f(x_i) - y_i)^2$, in which case the estimate can be shown to be:

$$\widehat{f}_n(x) = K(x, X)(K + \lambda I)^{-1} y,$$

where $K(x, X) = (K(x, x_1), \ldots, K(x, x_n)) \in \mathbb{R}^n$, and $K \in \mathbb{R}^{n \times n}$ is the so-called kernel gram matrix with $K_{ij} = K(x_i, x_j)$, and $y = (y_1, \ldots, y_n) \in \mathbb{R}^n$ is the set of responses. In the case where $\lambda = 0$, this yields the estimator: $\widehat{f}_n(x) = K(x, X)K^{-1}y$.

2 Random Features

Suppose the Mercer kernel has the spectral decomposition:

$$K(x, y) = \int \lambda(w) \phi(w, x) \phi(w, y) dw,$$

where $\{\phi(w, \cdot)\}$ are called the eigenfunctions of the kernel, and $\{\lambda(w)\}$ the corresponding eigenvalues. Then, an alternative characterization of the RKHS $\mathcal{H}_K$ is:

$$\mathcal{H}_K = \left\{ f(x) = \int \theta(w)\phi(w, x)dw \ \Big| \ \int \frac{\theta^2(w)}{\lambda(w)} < \infty \right\}.$$

A key advantage of such a spectral characterization is that the kernels, as well as the functions from the corresponding RKHS, can be approximated via a finite dimensional feature map. Suppose we draw m random samples $W := \{w_j\}_{j=1}^m$ iid from $p(w) := \lambda(w)/\int \lambda(w)dw$. This then specifies the following finite-dimensional feature map:

$$\Phi_W(x) = \left(\frac{1}{\sqrt{m}}\phi(w_1, x), \ldots, \frac{1}{\sqrt{m}}\phi(w_m, x) \right),$$

that is indexed by the random draw W. Thus, the feature map $\Phi_W(x)$ itself is random. In contrast to learning a function in the RKHS corresponding to K, one can then simply learn a linear function of these random features, an approach, which naturally, is called "random features" (Rahimi & Recht 2008, 2009). But would doing so approximate the RKHS, and would the Euclidean inner product of the random feature map approximate the kernel?

Approximating the kernel. Consider the empirical kernel $\widehat{K}(x, y) = \frac{1}{m}\sum_{j=1}^m \phi(w_j, x)\phi(w_j, y)$, which is simply the Euclidean dot product of the feature maps $\Phi_W(x)$ and $\Phi_W(y)$. Rahimi & Recht (2008) then showed that for any compact set $\mathcal{M}$,

$$\mathbb{P}\left[\sup_{(x,y)\in\mathcal{M}} (\widehat{K}(x, y) - K(x, y)) \geq \epsilon \right] \prec \frac{\mathrm{diam}^2(\mathcal{M})}{\epsilon^2} \exp(-m\,\epsilon^2).$$

Approximating the RKHS. Consider the sub-class $\mathcal{H}_C \subset \mathcal{H}_K$ given by:

$$\mathcal{H}_C = \left\{ \int \theta(w)\phi(w, x)dw \ \Big| \ |\theta(w)| < Cp(w) \right\},$$

for some constant $C > 0$. Consider its empirical counterpart:

$$\widehat{\mathcal{H}}_{W;C} = \left\{ f(x) = \sum_{j=1}^m \theta_j\phi(w_j, x) \ \Big| \ |\theta_j| \leq C/\sqrt{m} \right\}.$$

Consider the estimator

$$\widehat{f}_W \in \arg\inf_{f\in\widehat{\mathcal{H}}_{W;C}} \widehat{R}_n(f).$$

Rahimi & Recht (2009) then showed that with probability at least $1 - \delta$:

$$R(\widehat{f}_W) - \min_{f\in\mathcal{H}_C} R(f) \prec \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right) \sqrt{\log\frac{1}{\delta}}.$$

2

In particular, they showed that the approximation error scales as:

$$\inf_{f \in \widehat{\mathcal{H}}_{W;C}} \mathcal{R}(f) - \inf_{f \in \mathcal{H}_C} R(f) \prec \frac{1}{m},$$

and that the estimation error scales as:

$$\sup_{f \in \widehat{\mathcal{H}}_{W;C}} (\widehat{R}(f) - R(f)) \prec \frac{1}{n}.$$

Rahimi & Recht (2009) moreover showed that linear combinations of such randomly weighted features $\{\phi(w_j, x)\}$ often outperform approaches that also aim to learn these weights, such as in additive boosting.

Random Features as a Kernel Machine.   The previous results show that random features i.e. learning a linear parameterization of a random feature map drawn from the spectral decomposition of the kernel, would *approximate* a kernel machine (i.e. learning a function from the corresponding RKHS) with the kernel K. The approximation results hinge on the fact that random features are also a kernel machine, just with a "random features kernel" $\widehat{K}(x, y) = \langle \Phi_W(x), \Phi_W(y) \rangle$, which is an approximation of the kernel K.

3   Generalized Random Features and Top Layer Training of DNNs

Suppose we have a generic random feature map, $\Phi_W(X)$, that is indexed by a random vector W, but that is not necessarily derived via a spectral decomposition of a Mercer Kernel. Suppose we learn linear functions of this feature map, $f_\theta(x) = \sum_{j=1}^m \theta_j \Phi_{W;j}(x)$. We could then consider this as a kernel machine with the random kernel $\widehat{K}(x, y) = \langle \Phi_W(x), \Phi_W(y) \rangle$, which could be viewed as a randomized approximation of the kernel $K(x, y) = \mathbb{E}_W \langle \Phi_W(x), \Phi_W(y) \rangle$. This broadened perspective of random features can be used to connect DNNs and kernels: if we can approximate a DNN (or any similar complex model) via a linear function of a random feature map, then we could in turn connect it to kernel machines with some specific kernel.

The most natural approach to extract such a random feature map from a DNN is implicit in the common approach of freezing all but the top layer of some pre-trained DNN, and just fitting the top layer. Let $\Phi_W(x)$ denote the mapping from the inputs to the penultimate layer: this is random because the pre-trained DNN weights (in this case, of all but the last layer) depend on random training points, and some random initialization. They thus provide a generic random feature map, that in turn specify the random kernel $\widehat{K}$, with expectation K as detailed above. A simpler instance of this is where we simply randomly initialize all

but the top layer, and then just fit the top layer. The corresponding kernel $\widehat{K}$ in that case can be shown to converge to a specific compositional kernel that we will discuss in the next section.

3.1 Kernel Gradient Descent

You might have seen the phrase "kernel gradient descent" in some recent papers. Suppose we fix the random features $\Phi_W(\cdot)$ (for instance extracted from a randomly wired or pre-trained DNN as discussed above), and learn the linear functions $f_\theta(\cdot)$ given training data $\{(x_i, y_i)\}_{i=1}^n$ by solving for:

$$\inf_\theta \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f_\theta(x_i), y_i),$$

via gradient descent. The gradient descent path is then given by:

$$\frac{d\theta}{dt} = -\eta \frac{1}{n} \sum_i \frac{\partial \mathcal{L}(f_\theta(x_i), y_i)}{\partial f_\theta} \Phi_W(x_i),$$

where we use the shorthand

$$\Phi_W(x) = \left(\frac{1}{\sqrt{m}} \phi(w_1, x), \ldots, \frac{1}{\sqrt{m}} \phi(w_m, x) \right).$$

It thus follows that:

$$\frac{df_\theta(x)}{dt} = -\eta \frac{1}{n} \sum_i \frac{\partial \mathcal{L}(f_\theta(x_i), y_i)}{\partial f_\theta} \langle \Phi_W(x_i), \Phi_W(x) \rangle$$

$$= -\eta \frac{1}{n} \sum_i \frac{\partial \mathcal{L}(f_\theta(x_i), y_i)}{\partial f_\theta} \widehat{K}(x_i, x),$$

which has as its limit point precisely the unregularized RKHS regression estimate with the random features kernel $\widehat{K}$. For instance, for the squared loss $\mathcal{L}(f(x), y) = \frac{1}{2}(f(x) - y)^2$, the limit point is precisely $\widehat{K}(x, X)\widehat{K}^{-1}y$. Thus gradient descent on the linear coefficients of random features is also termed "kernel gradient descent" with the kernel $\widehat{K}$ in some recent papers.

4 Randomly Wired DNNs and Gaussian Processes

We have just seen that just training the top layer of a DNN, while freezing the rest of the layers to some random initialization, or some pre-trained random weights, can be connected to a corresponding kernel machine. But what if we do not train even the top layer, and

simply consider a DNN with randomly set weights. Would this still compute something useful? As it turns out, it still serves as a very useful Gaussian Process prior over possible functions (Lee et al. 2017).

Consider an L layer DNN, with n_ℓ units in the ℓ-th layer, for $\ell \in \{0, \ldots, L\}$, where we have $n_0 = d$ features in the input layer. Denote the pre-activation output of the ℓ-th layer as $z^{[\ell]}$, and the output of the ℓ-th layer as $x^{[\ell]} = \phi(z^{[\ell]})$. Thus, the pre-activation computation at the ℓ-th layer is given as

$$z_i^{[\ell]} = b^{[\ell]} + \sum_{j=1}^{n_{\ell-1}} W_{ij}^{[\ell]} x_j^{[\ell-1]},$$

where $b^{[\ell]}$ is the scalar bias term, and $W^{(\ell)} \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$. Denote the set of all parameters by θ, and the final DNN output as $f_\theta(x) = x^{[L]}(x)$.

Consider the random initialization of the parameters as: $W_{ij}^{[\ell]} \sim \mathcal{N}(0, \sigma_w/n_\ell)$, and $b^{[\ell]} \sim \mathcal{N}(0, \sigma_b)$. Since we have that $z_i^{[1]} = b^{[1]} + \sum_{j=1}^d W_{ij}^{[1]} x_j$, it follows that $z_i^{[1]}$ and $z_{i'}^{[1]}$ are independent for $i \neq i'$, and hence so are $x_i^{[1]}$ and $x_{i'}^{[1]}$. Consider the next layer: $z_i^{[2]} = b^{[2]} + \sum_{j=1}^{n_1} W_{ij}^{[2]} x_j^{[1]}$. In the limit of infinite width, with $n_1 \to \infty$, by the Central Limit Theorem, it follows that $z_i^{[2]}(x)$ is Gaussian distributed. Moreover, any collection $(z_i^{[2]}(x_1), \ldots, z_i^{[2]}(x_m))$ is also multivariate Gaussian distributed. It thus follows that $z_i^{[2]}(x)$ is a Gaussian Process with mean $\mu^2(x) = \mathbb{E}[z^{[2]}(x)] = 0$, and covariance:

$$K^{[2]}(x, x') = \mathbb{E}[z_i^{[2]}(x) z_i^{[2]}(x')] = \sigma_b^2 + \sigma_w^2 C(x, x'),$$

where $C(x, x') = \mathbb{E}[x_i^{[1]}(x) x_i^{[1]}(x')]$ is obtained by integrating out $W^{[1]}$ and $b^{[1]}$.

We can generalize this argument to all layers of a DNN by induction. Suppose that $z_i^{[\ell-1]}$ is a Gaussian Process and independent for every j. Then by the same argument as above, in the limit $n_\ell \to \infty$, it follows that $z_i^{[\ell]}(x)$ is a Gaussian Process with mean zero, and covariance $K^{[\ell]}(x, x') = \sigma_b^2 + \sigma_w^2 \mathbb{E}_{z_i^{[\ell-1]} \sim \mathrm{GP}(0, K^{[\ell-1]})}[\phi(z_i^{[\ell-1]}(x)), \phi(z_i^{[\ell-1]}(x'))]$.

Since the expectation is only over two Gaussian variables, it can be simplified as:

$$K^{[\ell]}(x, x') = \sigma_b^2 + \sigma_w^2 F_\phi(K^{[\ell-1]}(x, x), K^{[\ell-1]}(x, x'), K^{[\ell-1]}(x', x')),$$

for some function $F_\phi(\cdot)$.

Note that in the base case, $K^{[1]}(x, x') = \mathbb{E}[z_i^{[1]}(x) z_i^{[1]}(x')] = \sigma_b^2 + \sigma_w^2 \frac{x^T x'}{d}$, where d is the dimensionality of the inputs x.

It thus follows that at initialization, infinite width DNNs are Gaussian Processes with a specific "deep" or compositional kernel that is as specified by the compositional recurrence above.

5

Consider the Gaussian Process regression model $y = f(x) + \epsilon$, where $f \sim \mathrm{GP}(0, K^{[L]})$, and $\epsilon \sim N(0, \sigma_\epsilon^2)$. Then, given data $\{(x_i, y_i)\}_{i=1}^n$, as our Bayesian predictor at a test point x, we can simply output the posterior mean:

$$K(x, X)(K(X, X) + \sigma_\epsilon^2 I)^{-1} y,$$

where we have used the shorthand $K(x, X) = (K(x, x_1), \ldots, K(x, x_n))$.

5 Fully trained DNNs, Linearized Models

So far, we have seen that either just training the top layer, freezing the rest, or not training any layer, just using a fully randomly wired DNN, can both be connected to kernel machines. But what if we trained all the DNN layers. Could this still be connected to a kernel machine?

More generally, consider training a complex parametric function $f_\theta(x)$ (for instance a DNN) via gradient descent:
$$\frac{d\theta}{dt} = -\eta \frac{1}{n} \sum_i \frac{\partial \mathcal{L}(f_\theta(x_i), y_i)}{\partial f_\theta} \frac{\partial f_\theta(x_i)}{\partial \theta},$$

so that it follows that:

$$\frac{df_\theta(x)}{dt} = -\eta \frac{1}{n} \sum_i \frac{\partial \mathcal{L}(f_\theta(x_i), y_i)}{\partial f_\theta} \left\langle \frac{\partial f_\theta(x_i)}{\partial \theta}, \frac{\partial f_\theta(x)}{\partial \theta} \right\rangle$$

$$= -\eta \frac{1}{n} \sum_i \frac{\partial \mathcal{L}(f_\theta(x_i), y_i)}{\partial f_\theta} \widehat{K}_t(x_i, x),$$

where

$$\widehat{K}_t(x, y) = \left\langle \frac{\partial f_\theta(x)}{\partial \theta}, \frac{\partial f_\theta(y)}{\partial \theta} \right\rangle.$$

When f_θ is a DNN, Jacot et al. (2018) termed the corresponding kernel $\widehat{K}_t$ as the Neural Tangent Kernel at the current iterate θ_t. Jacot et al. (2018) moreover showed the following two results: (a) as the widths of the layers sequentially approach infinity (so-called infinitely wide DNNs), the random NTKs at the initialization $\widehat{K}_0$ converge to a deterministic compositional kernel, which can be derived via a recurrence similar to that of the Gaussian Process kernel above; and (b) uniformly over $t \in [0, T]$ for some finite time horizon $T > 0$, the random NTKs at the iterates θ_t also converge to that same limit determinsitic NTK.

Arora et al. (2019) further refined this analysis providing non-asymoptotic deviation bounds. Specifically, they showed that:

Theorem 1 $\forall \epsilon > 0, \delta \in (0, 1)$, *suppose that* $\min_{\ell \in [L]} n_\ell \prec \log(1/\delta)/\epsilon^4$. *Then for any inputs* $x, x' \in \mathbb{R}^d$, *such that* $\|x\| \leq 1, \|x'\| \leq 1$, *with probability at least* $1 - \delta$, *there exists a*

deterministic limit kernel K_∞ such that:

$$|\widehat{K}_t(x, x') - K_\infty(x, x')| \prec \epsilon.$$

5.1 Linearization

The results above essentially show that in infinite width limit, DNNs are essentially behave similar to the generalized random features machines as in the previous section: linear machines of some random feature maps, which in turn specify some random kernels, which are close to their deterministic expectations. But why do they have such linear behavior?

Consider a complex parametric function f_θ (for instance a DNN). When this function is highly overparameterized, and we take relatively fewer descent steps, then it is likely that each individual parameter moves only slightly away from its initial value, to have the overall function value with its hundreds of millions of parameters to move sufficiently far from its initial value, and in particular to fit the data well. In such cases, where each individual parameter only moves a small amount even with full training, we can approximate the learning of f_θ to simply learning the coefficients of a linearization (Lee et al. 2019):

$$f_w^{\text{LIN}}(x) = f_{\theta_0}(x) + \left\langle w, \frac{\partial f_\theta(x)}{\partial \theta}\Big|_{\theta=\theta_0} \right\rangle.$$

Denote $\phi_{\theta_0}(x) = \frac{\partial f_\theta(x)}{\partial \theta}\big|_{\theta=\theta_0}$. Suppose we initialize θ_0 randomly, for instance via iid standard Gaussians: $[\theta_0]_i \sim N(0, 1)$. It can then be seen that training this linear model is essentially random features regression corresponding to the kernel:

$$K_{\text{NTK}}(x, x') = \text{Cov}_{\theta_0}[\phi_{\theta_0}(x), \phi_{\theta_0}(x')].$$

In the case of DNNs, this can be seen to be precisely the Neural Tangent Kernel (NTK) at initialization.

5.2 Lazy Training

Chizat et al. (2019) recently argued that such linearization (or what they call "lazy training") is not due to overparameterization per se, but rather due to a particular scaling of the parameters relative to the gradient descent step size. Their high level insight is that if the step size is too small relative to the scale of the parameters, we might naturally expect to be in the "lazy training" or linearized regime where the parameters inividually do not move too far from the initial parameters.

Towards formalizing this, consider the rescaled risk function:

$$R_\alpha(\theta) = \frac{1}{\alpha^2} R(\alpha f_\theta),$$

as well as its linearization:

$$R_\alpha^{\mathrm{LIN}}(\theta) = \frac{1}{\alpha^2} R(\alpha f_\theta^{\mathrm{LIN}}),$$

where

$$f_\theta^{\mathrm{LIN}}(x) = f_{\theta_0}(x) + \left\langle \theta, \frac{\partial f_\theta(x)}{\partial \theta}\Big|_{\theta=\theta_0} \right\rangle,$$

is the linearization of f_θ around some initial parameter θ_0. Now consider both gradient descent on the full rescaled risk:

$$\frac{d\theta_\alpha}{dt} = -\nabla R_\alpha(\theta),$$

as well as gradient descent on the linearized rescaled risk:

$$\frac{d\theta_\alpha^{\mathrm{LIN}}}{dt} = -\nabla R_\alpha^{\mathrm{LIN}}(\theta).$$

Chizat et al. (2019) then showed that given a finite time horizon $T > 0$:

$$\sup_{t\in[0,T]} \|\theta_\alpha(t) - \theta_0\| = O(1/\alpha) \tag{1}$$

$$\sup_{t\in[0,T]} \|\theta_\alpha(t) - \theta^{\mathrm{LIN}}(t)\| = O(1/\alpha^2) \tag{2}$$

so that the gradient descent iterates remain close to the initial parameter, and the gradient descent path for the linearized function remains close to the path for the original function itself.

6 The representational power of (compositional) kernel methods

The previous sections showed that at least in some over-parameterized or parameter scaling regimes, learning complex functions such as DNNs is essentially similar to kernel machines for some fixed, if a deep or compositionally specified kernel function. But how good are such kernel machines? Do we actually get the flexibility of arbitrary DNNs?

Chizat et al. (2019) empirically showed that lazy training (such as in learning linearized models in overparameterized settings) is much worse empirically than "non-lazy" training where the parameters are no longer in the "lazy training" regimes. Ghorbani et al. (2019) bolstered this with a very powerful set of theoretical results that rigorously analyzed the approximation and generalization errors of such kernel machines, ranging over both random as well as deterministic kernel machines discussed above.

Specifically, they considered the setting where the inputs $x_i \sim \text{Unif}(S^{d-1}(\sqrt{d}))$, where $S^{d-1}(r)$ denotes the sphere with radius r in d dimensions. And the responses $y_i = f^*(x_i)$, for some arbitrary $f^* : S^{d-1}(\sqrt{d}) \mapsto \mathbb{R}$.

For the random kernel machines, let N be the dimensionality of finite-dimensional random feature maps. They then showed that if $d^{\ell+\delta} \leq N \leq d^{\ell+1-\delta}$, then the shallow random kernel machines have the same approximation error as fitting a linear function over all monomials of degree ℓ in x, while the deeper Neural Tangent Kernel based random kernel machines have the same approximation error as fitting a linear function over all monomials of degree $\ell + 1$ in x. They also considered general deterministic kernels that are rotationally invariant over the sphere $S^{d-1}(\sqrt{d})$, which might have zero or negligible approximation error. But in such cases, the generalization error is in turn large. Let n denote the number of samples. They then showed that when $d^{\ell+\delta} \leq n \leq d^{\ell+1-\delta}$, then the generalization error of these kernel machines scales as the approximation error of linear regression over all monomials of degree ℓ in x.

They also coupled this with a very simple separation result to show that this does not entail that DNNs itself have large approximation or generalization error. Specifically, consider the case where the target function $f^* = \sigma(w^* \cdot x)$ is a single neuron. In such a case, under the distributional assumptions on the training inputs above, it can be shown that gradient descent over the empirical mean squared error recovers the true parameter w^*. But using say the NTK kernel machine with $d^{\ell+\delta} \leq N \leq d^{\ell+1-\delta}$ would have approximation error of linear regression over all monomials of degree ℓ in x, which could be large if σ cannot be approximated well by polynomials of degree atmost ℓ.

References

Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R. & Wang, R. (2019), On exact computation with an infinitely wide neural net, *in* 'Advances in Neural Information Processing Systems', pp. 8139–8148.

Chizat, L., Oyallon, E. & Bach, F. (2019), On lazy training in differentiable programming, *in* 'Advances in Neural Information Processing Systems', pp. 2933–2943.

Ghorbani, B., Mei, S., Misiakiewicz, T. & Montanari, A. (2019), 'Linearized two-layers neural networks in high dimension', *arXiv preprint arXiv:1904.12191* .

Jacot, A., Gabriel, F. & Hongler, C. (2018), Neural tangent kernel: Convergence and generalization in neural networks, *in* 'Advances in neural information processing systems', pp. 8571–8580.

Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J. & Sohl-Dickstein, J. (2017), 'Deep neural networks as gaussian processes', *arXiv preprint arXiv:1711.00165* .

Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J. & Pennington, J. (2019), Wide neural networks of any depth evolve as linear models under gradient descent, *in* 'Advances in neural information processing systems', pp. 8570–8581.

Rahimi, A. & Recht, B. (2008), Random features for large-scale kernel machines, *in* 'Advances in neural information processing systems', pp. 1177–1184.

Rahimi, A. & Recht, B. (2009), Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning, *in* 'Advances in neural information processing systems', pp. 1313–1320.