Adapting Kernel Representations Online Using Submodular Maximization

Matthew Schlegel [1] **Yangchen Pan** [1] **Jiecao Chen** [1] **Martha White** [1]

Abstract

Kernel representations provide a nonlinear representation, through similarities to prototypes, but require only simple linear learning algorithms given those prototypes. In a continual learning setting, with a constant stream of observations, it is critical to have an efficient mechanism for sub-selecting prototypes amongst observations. In this work, we develop an approximately submodular criterion for this setting, and an efficient online greedy submodular maximization algorithm for optimizing the criterion. We extend streaming submodular maximization algorithms to continual learning, by removing the need for multiple passes—which is infeasible—and instead introducing the idea of coverage time. We propose a general block-diagonal approximation for the greedy update with our criterion, that enables updates linear in the number of prototypes. We empirically demonstrate the effectiveness of this approximation, in terms of approximation quality, significant runtime improvements, and effective prediction performance.

1. Introduction

Kernel representations provide an attractive approach to representation learning, by facilitating simple linear prediction algorithms and providing an interpretable representation. A kernel representation consists of mapping an input observation into similarity features, with similarities to a set of prototypes. Consequently, for an input observation, a prediction can be attributed to those prototypes that are most similar to the observation. Further, the transformation to similarity features is non-linear, enabling nonlinear function approximation while using linear learning algorithms that simply optimize for weights on these transformed features. Kernel representations are universal func-

tion approximators[1] and the flexibility in choosing the kernel (similarity function) has enabled impressive prediction performance for a range of settings, including speech (Huang et al., 2014), computer vision (Mairal et al., 2014), and object recognition (Lu et al., 2014).

In a continual learning setting, such as in online learning or reinforcement learning, there is a constant, effectively unending stream of data, necessitating some care when using kernel representations. The issue arises from the choice of prototypes. Before the advent of huge increases in dataset sizes, a common choice was to use all of the training data as prototypes. This choice comes from the representer theorem, which states that for a broad class of functions, the empirical risk minimizer is a linear weighting of similarity features, to a set of prototypes that consists of the training data. For continual learning, however, the update should be independent of the the total number of samples— which is not clearly defined for continual learning. Conversely, we want to permit selection of a sufficiently large number of prototypes, to maintain sufficient modeling power. For efficient, continual updating, therefore, we require per-step prototype selection strategies that are approximately linear in the number of prototypes.

Currently, most algorithms do not satisfy the criteria for a continual learning setting. Incremental selection of prototypes has been tackled in a wide range of areas, due to the fundamental nature of this problem. Within the streaming community, approaches typically assume that the batch of data, though large, is accessible and fixed. The most related of these areas[2] include active set selection for Gaussian process regression (Seeger et al., 2003), with streaming submodular maximization approaches (Krause et al., 2008b;a; Badanidiyuru et al., 2014); incremental Nystrom methods within kernel recursive least-squares (KRLS)

---

[1]Department of Computer Science, Indiana University, Bloomington. Correspondence to: Martha White <martha@indiana.edu>.

[1]Radial basis function networks are an example of a kernel representation, that have been shown to be universal function approximators (Park & Sandberg, 1991). Further, the representer theorem further characterizes the approximation capabilities under empirical risk minimization for a broad class of functions.

[2]Facility location, k-medians and k-centers are three problems that focus on selecting representative instances from a set (c.f. (Guha et al., 2003)). The criteria and algorithms are not designed with the intention to use the instances for prediction and so we do not consider them further here.

(Rudi et al., 2015)[3]; and functional gradients that sample random bases which avoid storing prototypes but require storing n scalars, for n training samples (Dai et al., 2014).

Kernel representation algorithms designed specifically for the online setting, on the other hand, are typically too computationally expensive in terms of the number of prototypes. Kernel least-mean squares (KLMS) algorithms use stochastic updates, maintaining the most recent prototypes and truncating coefficients on the oldest (Kivinen et al., 2010; Schraudolph et al., 2006; Cheng et al., 2007); though efficient given sufficient truncation, this truncation can introduce significant errors (Van Vaerenbergh & Santamaria, 2013). Random feature approximations (Rahimi & Recht, 2007) can be used online, but require a significant number of random features. Gaussian process regression approaches have online variants (Csató & Opper, 2006; Cheng & Boots, 2016), however, they inherently require at least quadratic computation to update the variance parameters. KRLS can be applied online, but has a threshold parameter that makes it difficult to control the number of prototypes and requires quadratic computation and space (Engel et al., 2004). More efficient coherence heuristic have been proposed (Richard et al., 2009; Van Vaerenbergh et al., 2010; Chen et al., 2013; Van Vaerenbergh & Santamaria, 2013), but provide no approximation quality guarantees.

In this work, we provide a simple and efficient greedy algorithm for selecting prototypes for continual learning, by extending recent work in prototype selection with submodular maximization. We introduce a generalized coherence criterion for selecting prototypes, which unifies two previously proposed criteria: the coherence criterion and the log determinant. Because this criterion is (approximately) submodular, we pursue a generalization to streaming submodular maximization algorithms. We avoid the need for multiple passes over the data—which is not possible in continual learning— by introducing the idea of coverage time, which reflects that areas of the observation space are repeatedly visited under sufficient mixing. We prove that our online submodular maximization achieves an approximation-ratio of 1/2, with a small additional approximation introduced due to coverage time and from using an estimate of the submodular function. We then provide a linear-time algorithm for approximating one instance of our submodular criterion, by exploiting the block-diagonal form of the kernel matrix. We empirically demonstrate that this approximation closely matches the true value, despite using significantly less computation, and show effective prediction performance using the corresponding kernel representation.

---

[3]There is a large literature on fast Nystrom methods using related approaches, such as determinant point processes for sampling landmark points (Li et al., 2016). The primary goal for these methods, however, is to approximate the full kernel matrix.

2. Using kernel representations

A kernel representation is a transformation of observations into similarity features, consisting of similarities to prototypes. A canonical example of such a representation is a radial basis function network, with radial basis kernels such as the Gaussian kernel; however, more generally any kernel similarity can be chosen. More formally, for observations $\mathbf{x} \in \mathcal{X}$, the kernel representation consists of similarities to a set of prototypes $S = \{\mathbf{z}_1, \ldots, \mathbf{z}_b\} \subset \mathcal{X}$

$$\mathbf{x} \to [k(\mathbf{x}, \mathbf{z}_1), \ldots, k(\mathbf{x}, \mathbf{z}_b)] \in \mathbb{R}^b.$$

for kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$. The observations need not be numerical; as long as a similarity k can be defined between two observations, kernel representations can be used and conveniently provide a numeric feature vector in $\mathbb{R}^b$. We use the term prototype, instead of center, to emphasize that the chosen observations are representative instances, that are sub-selected from observed data.

A fundamental result for kernel representations is the representer theorem, with significant recent generalizations (Argyriou & Dinuzzo, 2014), which states that for a broad class of function spaces $\mathcal{H}$, the empirical risk minimizer $f \in \mathcal{H}$ on a training set $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ has the simple form

$$f(\cdot) = \sum_{i=1}^n \alpha_i k(\cdot, \mathbf{x}_i).$$

This result makes use of a key property: the kernel function can be expressed as an inner product, $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$ for some implicit expansion ϕ. The function f can be written $f = \sum_{i=1}^n \alpha_i \phi(\mathbf{x}_j)$, with

$$f(\mathbf{x}) = \langle \phi(\mathbf{x}), \sum_{i=1}^n \alpha_i \phi(\mathbf{x}_j) \rangle = \sum_{i=1}^n \alpha_i \langle \phi(\mathbf{x}), \phi(\mathbf{x}_j) \rangle.$$

It is typically impractical to use all $\mathbf{x}_i$ as prototypes, and a subset needs to be chosen. Recently, there has been several papers (Krause et al., 2008b;a; Krause & Gomes, 2010; Badanidiyuru et al., 2014) showing that prototypes can be effectively selected in the streaming setting using greedy submodular maximization on the log-determinant of the kernel matrix, $\mathbf{K}_S \in \mathbb{R}^{b \times b}$ where $(\mathbf{K}_S)_{ij} = k(\mathbf{z}_i, \mathbf{z}_j)$. Given some ground set Ω and its powerset $\mathcal{P}(\Omega)$, submodular functions $g : \mathcal{P}(\Omega) \to \mathbb{R}$ are set functions, with a diminishing returns property: the addition of a point to a given set increases the value less or equal than adding a point to a subset of that set. For prototype selection, the ground set considered are sets of all observations $\mathcal{X}$, so $S \subset \Omega = \mathcal{X}$. The log-determinant of the resulting kernel matrix, $\log \det \mathbf{K}_S$, is a submodular function of S. Though maximizing submodular functions is NP-hard, greedy approximation algorithms have been shown to obtain reasonable approximation ratios, even for the streaming setting

(Krause & Gomes, 2010; Badanidiyuru et al., 2014), and the resulting algorithms are elegantly simple and theoretically sound.

In the following sections, we derive a novel criterion for prototype selection, that includes the log-determinant as a special case. Then, we provide an efficient prototype selection algorithm for the continual learning setting, using submodular maximization.

3. Selecting kernel prototypes

Many prototype selection strategies are derived based on diversity measures. The coherence criterion (Engel et al., 2004) approximates how effectively the set of prototypes spans the set of given observations. The log-determinant measures the spread of eigenvalues for the kernel matrix, and is related to information gain (Seeger, 2004). These selection criteria are designed for a finite set of observed points; here, we step back and reconsider a suitable objective for prototype selection for continual learning.

Our goal is to select prototypes that minimize distance to the optimal function. In this section, we begin from this objective and demonstrate that the coherence criterion and log-determinant are actually upper bounds on this objective, and special cases of a more general such upper bound. The analysis justifies that the log-determinant is a more suitable criteria for continual learning, which we then pursue in the remainder of this work.

3.1. Criteria to select prototypes from a finite set

Let $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ be a set of points, with corresponding labels $\mathbf{y}_1, \ldots, \mathbf{y}_n$. We will not assume that this is a batch of data, but could rather consist of all possible observations for a finite observation space. Ideally, we would learn $S = \{\mathbf{z}_1, \ldots, \mathbf{z}_b\} \subset \mathcal{X}$ and corresponding $f_{S,\beta}(\cdot) = \sum_{i=1}^{b} \beta_j k(\cdot, \mathbf{x}_i)$ according to the loss

$$\min_{S \subset \mathcal{X}} \min_{\beta \in \mathbb{R}^b} \|f - f_{S,\beta}\|^2 \tag{1}$$

$$\text{where} \|f - f_{S,\beta}\| = \left\| \sum_{i=1}^{n} \alpha_i \phi(\mathbf{x}_i) - \sum_{j=1}^{b} \beta_j \phi(\mathbf{z}_j) \right\|$$

for the optimal f for the set of points $\mathcal{X}$. Because we do not have f, we derive an upper bound on this value. Introducing dummy variables $\beta_j^{(i)} \in \mathbb{R}$ such that $\beta_j = \sum_{i=1}^{n} \beta_j^{(i)}$,

$$(1) \leq \min_{S \subset \mathcal{X}, \beta \in \mathbb{R}^b} \sum_{i=1}^{n} \left\| \alpha_i \phi(\mathbf{x}_i) - \sum_{j=1}^{b} \beta_j^{(i)} \phi(\mathbf{z}_j) \right\|^2$$

$$= \min_{S \subset \mathcal{X}} \sum_{i=1}^{n} \min_{\beta_1^{(i)}, \ldots, \beta_b^{(i)}} \left\| \alpha_i \phi(\mathbf{x}_i) - \sum_{j=1}^{b} \beta_j^{(i)} \phi(\mathbf{z}_j) \right\|^2 \tag{2}$$

For $\mathbf{k}_i \doteq [k(\mathbf{x}_i, \mathbf{z}_1), \ldots, k(\mathbf{x}_i, \mathbf{z}_b)]$, the interior minimization can be re-written as

$$\min_{\beta^{(i)}} \left\| \alpha_i \phi(\mathbf{x}_i) - \sum_{j=1}^{b} \beta_j^{(i)} \phi(\mathbf{z}_j) \right\|^2$$

$$= \min_{\beta} \beta^{(i)\top} \mathbf{K}_S \beta - 2\alpha_i \beta^{(i)\top} \mathbf{k}_i + \alpha_i^2 k(\mathbf{x}_i, \mathbf{x}_i)$$

To provide more stable approximations, we regularize

$$(2) \leq \min_{S \subset \mathcal{X}} \sum_{i=1}^{n} \min_{\beta^{(i)}} \left\| \alpha_i \phi(\mathbf{x}_i) - \sum_{j=1}^{b} \beta_j^{(i)} \phi(\mathbf{z}_j) \right\|^2 + \lambda \|\beta^{(i)}\|_2^2$$

$$= \min_{S \subset \mathcal{X}} \sum_{i=1}^{n} \min_{\beta^{(i)}} \beta^{(i)\top} (\mathbf{K}_S + \lambda \mathbf{I})\beta - 2\alpha_i \beta^{(i)\top} \mathbf{k}_i$$

$$+ \alpha_i^2 k(\mathbf{x}_i, \mathbf{x}_i).$$

For $\lambda = 0$, the inequality is equality. Otherwise adding regularization theoretically increases the upper bound, though in practice will be key for stability.

Solving gives $\beta^{(i)} = \alpha_i (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i$, and so

$$\beta^{(i)\top} (\mathbf{K}_S + \lambda \mathbf{I})\beta - 2\alpha_i \beta^{(i)\top} \mathbf{k}_i + \alpha_i^2 k(\mathbf{x}_i, \mathbf{x}_i)$$

$$= \alpha_i^2 \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i - 2\alpha_i^2 \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i$$

$$+ \alpha_i^2 k(\mathbf{x}_i, \mathbf{x}_i)$$

$$= \alpha_i^2 k(\mathbf{x}_i, \mathbf{x}_i) - \alpha_i^2 \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i$$

We can now simplify the above upper bound

$$(2) \leq \min_{S \subset \mathcal{X}} \sum_{i=1}^{n} \left(\alpha_i^2 k(\mathbf{x}_i, \mathbf{x}_i) - \alpha_i^2 \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i \right)$$

and obtain equivalent optimization

$$\operatorname*{argmax}_{S \subset \mathcal{X}} \sum_{i=1}^{n} \alpha_i^2 \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i.$$

This criteria closely resembles the coherence criterion (Engel et al., 2004). The key idea for the coherence criterion is to add a prototype $\mathbf{x}_i$, to kernel matrix $\mathbf{K}_S$ if $1 - \mathbf{k}_i \mathbf{K}_S^{-1} \mathbf{k}_i \leq \nu$ for some threshold parameter ν. The coherence criterion,

$$\operatorname*{argmax}_{S \subset \mathcal{X}} \sum_{i=1}^{n} \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i$$

therefore, can be seen as an upper bound on the distance the optimal function, with further relaxation because $\sum_{i=1}^{n} \alpha_i^2 \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i \leq \max\{\alpha_1^2, \ldots, \alpha_n^2\} \sum_{i=1}^{n} \mathbf{k}_i^\top (\mathbf{K}_S + \lambda \mathbf{I})^{-1} \mathbf{k}_i$.

The relationship to another popular criterion—the log determinant— arises when we consider the extension to an infinite state space.

3.2. Criteria to select prototypes from an infinite set

The criterion above can be extended to an uncountably infinite observation space $\mathcal{X}$. For this setting, the optimal $f = \int_{\mathcal{X}} \omega(\mathbf{x})\phi(\mathbf{x})d\mathbf{x}$, for a function $\omega : \mathbb{R}^d \rightarrow \mathbb{R}$. Let $\mathbf{k}(\mathbf{x}, S) = [k(\mathbf{x}, \mathbf{z}_1), \ldots, k(\mathbf{x}, \mathbf{z}_b)]$ Then, using a similar analysis to above,

$$\min_{S \subset \mathcal{X}} \min_{\boldsymbol{\beta} \in \mathbb{R}^b} \|f - f_{S,\beta}\|^2 \leq \min_{S \subset \mathcal{X}} \int_{\mathcal{X}} \omega(\mathbf{x})^2 k(\mathbf{x}, \mathbf{x})d\mathbf{x}$$
$$- \int_{\mathcal{X}} \omega(\mathbf{x})^2 \mathbf{k}(\mathbf{x}, S)^\top (\mathbf{K}_S + \lambda\mathbf{I})^{-1}\mathbf{k}(\mathbf{x}, S)d\mathbf{x}.$$

and so the resulting goal is to optimize

$$\underset{S \subset \mathcal{X}}{\operatorname{argmax}} \int_{\mathcal{X}} \omega(\mathbf{x})^2 \mathbf{k}(\mathbf{x}, S)^\top (\mathbf{K}_S + \lambda\mathbf{I})^{-1}\mathbf{k}(\mathbf{x}, S)d\mathbf{x}. \quad (3)$$

This provides a nice relation to the log determinant criterion, with normalized kernels[4]: $k(\mathbf{z}, \mathbf{z}) = 1$. If $\mathbf{k}(\mathbf{x}, S)$ maps to a unique kernel vector $\mathbf{k} \in [0, 1]^b$, and the function $\mathbf{k}(\cdot, S)$ also maps onto $[0, 1]^b$, then for $\omega(\mathbf{x}) = 1$,

$$\int_{\mathcal{X}} \omega(\mathbf{x})^2 \mathbf{k}(\mathbf{x}, S)^\top (\mathbf{K}_S + \lambda\mathbf{I})^{-1}\mathbf{k}(\mathbf{x}, S)d\mathbf{x}$$
$$= \int \mathbf{k}^\top (\mathbf{K}_S + \lambda\mathbf{I})^{-1}\mathbf{k} \, d\mathbf{k}$$
$$= \det(\mathbf{K}_S + \lambda\mathbf{I}).$$

In general, it is unlikely to have a bijection $\mathbf{k}(\cdot, S)$. More generally, we can obtain the above criterion by setting the coefficient function ω so that each possible kernel vector $\mathbf{k} \in \mathbb{R}^b$ has uniform weighting, or implicitly so the integration is uniformly over $\mathbf{k} \in [0, 1]^b$. Because $\log$ is monotonically increasing, maximizing $\det(\mathbf{K}_S + \lambda\mathbf{I})$ with a fixed b is equivalent to maximizing $\log \det(\mathbf{K}_S + \lambda\mathbf{I})$.

This derivation of an upper bound on the distance to the optimal function provides new insights into the properties of the log-determinant, clarifies the connection between the coherence criterion and the log-determinant, and suggesting potential routes for providing criteria based on the prediction utility of a prototype. The choice of weighting ω to obtain the log-determinant removes all information about the utility of a prototype and essentially assumes a uniform distribution over the kernel vectors $\mathbf{k}$. For more general coefficient functions ω, let $\boldsymbol{\mu}(S) = \mathbb{E}[\mathbf{k}(\boldsymbol{X}, S)]$ and $\boldsymbol{\Sigma}(S) = \operatorname{Cov}(\mathbf{k}(\boldsymbol{X}, S))$, where the expectations are according to density ω^2/c for normalizer $c = \int_{\mathcal{X}} \omega(\mathbf{x})d\mathbf{x}$. By the quadratic expectations properties (Brookes, 2004)

$$(3) = \underset{S \subset \mathcal{X}}{\operatorname{argmax}} \operatorname{tr}((\mathbf{K}_S + \lambda\mathbf{I})^{-1}\boldsymbol{\Sigma}(S))$$
$$+ \boldsymbol{\mu}(S)^\top (\mathbf{K}_S + \lambda\mathbf{I})^{-1}\boldsymbol{\mu}(S). \quad (4)$$

---

[4]A kernel can be normalized by $k(\mathbf{x}, \mathbf{z})/\sqrt{k(\mathbf{z}, \mathbf{z})k(\mathbf{x}, \mathbf{x})}$.

This more general form in (4) enables prototypes to be more highly weighted based on the magnitude of values in ω. We focus in this work first on online prototype selection for the popular log-determinant, and leave further investigation into this more general criteria to future work. We nonetheless introduce the form here to better motivate the log-determinant, as well as demonstrate that the above analysis is amenable to a host of potential directions for more directed prototype selection.

4. Online submodular maximization

In this section, we introduce an OnlineGreedy algorithm for submodular maximization, to enable optimization of the prototype selection objective from an online stream of data. Current submodular maximization algorithms are designed for the streaming setting, which deals with incrementally processing large but fixed datasets. Consequently, the objectives are specified for a finite batch of observations and the algorithms can do multiple passes over the dataset. For the online setting, both of these conditions are restrictive. We show that, with a minor modification to StreamGreedy (Krause & Gomes, 2010), we can obtain a comparable approximation guarantee that applies to the online setting.

We would like to note that there is one streaming algorithm, called Sieve Streaming, designed to only do one pass of the data and avoid too many calls to the submodular function (Badanidiyuru et al., 2014); however, it requires keeping parallel solutions, which introduces significant complexity and which we found prohibitively expensive. In our experiments, we show it is significantly slower than our approach and found it typically maintained at least 500 parallel solutions. For this reason, we opt to extend the simpler StreamGreedy algorithm, and focus on efficient estimates of the submodular function, since we will require more calls to this function than Sieve Streaming.

Our goal is to solve the submodular maximization problem

$$\max_{S \subset \mathcal{X}:|S| \leq b} g(S) \quad (5)$$

where $\mathcal{X}$ is a general space of observations and g is a submodular function. The key modification is to enable $\mathcal{X}$ to be a large, infinite or even uncountable space. For such $\mathcal{X}$, we will be unable to see all observations, let alone make multiple passes. Instead, we will use a related notion to mixing time, where we see a cover of the space.

The greedy algorithm consists of greedily adding in a new prototype if it is an improvement on a previous prototype. The resulting greedy algorithm—given in Algorithm 1— is similar to StreamGreedy, and so we term it OnlineGreedy. The algorithm queries the submodular function on each set, with a previous prototype removed and the new observation added. To make this efficient, we will rely on us-

---

Algorithm 1 OnlineGreedy

Input: threshold parameter ϵ_t, where a prototype is only added if there is sufficient improvement

$S_0 \leftarrow \emptyset$

for $t = 1 : b$ **do** $S_t \leftarrow S_{t-1} \cup \{\mathbf{x}_t\}$

while interacting, $t = b + 1, \ldots$ **do**

$\quad \mathbf{z}' = \underset{\mathbf{z} \in S_{t-1}}{\mathrm{argmax}} \, \hat{g}(S_{t-1} \backslash \{\mathbf{z}\} \cup \{\mathbf{x}_t\})$

$\quad S_t \leftarrow S_{t-1} \backslash \{\mathbf{z}'\} \cup \{\mathbf{x}_t\}$

$\quad$ **if** $\hat{g}(S_t) - \hat{g}(S_{t-1}) < \epsilon_t$ **then**

$\qquad S_t \leftarrow S_{t-1}$

---

ing only an approximation to the submodular function g. We will provide a linear-time algorithm—in the number of prototypes— for querying replacement to all prototypes, as opposed to a naive solution which would be cubic in the number of prototypes. This will enable us to use this simple greedy approach, rather than more complex streaming submodular maximization approaches that attempt to reduce the number of calls to the submodular function.

We bound approximation error, relative to the optimal solution. We extend an algorithm that uses multiple passes; our approach suggests more generally how algorithms from the streaming setting can be extended to an online setting. To focus the on this extension, we only consider submodular functions here; in Appendix B, we generalize the result to approximately submodular functions. Many set functions are approximately submodular, rather than submodular, but still enjoy similar approximation properties. The log-determinant is submodular, however, it is more likely that, for the variety of choices for ω, the generalized coherence criterion is only approximately submodular. For this reason, we provide this generalization to approximate submodularity, as it further justifies the design of (approximately) submodular criteria for prototype selection.

We compare our solution to the optimal solution

$$S^* = \underset{S \subset \mathcal{X} : |S| \leq b}{\mathrm{argmax}} \, g(S) = \{\mathbf{z}_1^*, \ldots, \mathbf{z}_b^*\}.$$

Assumption 1 (Submodularity). *g is monotone increasing and submodular.*

Assumption 2 (Approximation error). *We have access to a set function $\hat{g}$ that approximates g: for some $\epsilon_f \geq 0$ for all $S \subset \mathcal{X}$, with $|S| \leq b$,*

$$|\hat{g}(S) - g(S)| \leq \epsilon_f$$

Assumption 3 (Submodular coverage time). *For a fixed $\epsilon_r > 0$ and $\delta > 0$ there exists a $\rho \in \mathbb{N}$ such that for all $S \subset \mathcal{X}$ where $|S| \leq b$, with probability $1 - \delta$, for any $\mathbf{z}^* \in S^*$ an observation $\mathbf{x}$ is observed within ρ steps (starting from any point in $\mathcal{X}$) that is similar to $\mathbf{z}^*$ in that*

$$|g(S \cup \{\mathbf{x}\}) - g(S \cup \{\mathbf{z}^*\})| \leq \epsilon_r.$$

This final assumption characterizes that the environment is sufficiently mixing, to see a cover of the space. We introduce the term coverage, instead of cover time for finite-state, to indicate a relaxed notion of observing a covering of the space rather than observing all states.

For simplicity of the proof, we characterize the coverage time in terms of the submodular function. We show that the submodular function we consider—the log-determinant—satisfies this assumption, given a more intuitive assumption that instead requires that observations be similar according to the kernel. The statement and proof are in Appendix A.

Now we prove our main result.

Theorem 1. *Assume Assumptions 1-3 and that $g(S^*)$ is finite and $g(\emptyset) \geq 0$. Then, for $t > \rho g(S^*)/\epsilon_t$, all sets S_t chosen by OnlineGreedy using $\hat{g}$ satisfy, with probability $1 - \delta$,*

$$g(S_t) \geq \frac{1}{2} g(S^*) - \frac{b}{2}(\epsilon_r + 2\epsilon_f + \epsilon_t)$$

Proof: The proof follows closely to the proof of Krause & Gomes (2010, Theorem 4). The key difference is that we cannot do multiple passes through a fixed dataset, and instead use submodular coverage time.

Case 1: There have been $t \geq \rho g(S^*)/\epsilon_t$ iterations, and S_t has always changed within ρ iterations (i.e., there has never been ρ consecutive iterations where S_t remained the same). This mean that for each ρ iterations, $\hat{g}(S_t)$ must have been improved by at least ϵ_t, which is the minimum threshold for improvement. This means that over the t iterations, $\hat{g}(S_0)$ has improved by at least ϵ_t each ρ,

$$\hat{g}(S_0) + \epsilon_t t/\rho \geq \epsilon_t t/\rho = \hat{g}(S^*) \geq g(S^*) - \epsilon_f$$

The solution is within ϵ_f of $g(S^*)$, and we are done.

Case 2: At some time t, S_t was not changed for ρ iterations, i.e., $S_{t-\rho} = S_{t-\rho-1} = \ldots S_t$. Order the prototypes in the set as $\mathbf{z}_i = \mathrm{argmax}_{\mathbf{z} \in S_t} g(\{\mathbf{z}_1, \ldots, \mathbf{z}_{i-1}\} \cup \{\mathbf{z}\})$, with

$$\delta_i = g(\{\mathbf{z}_1, \ldots, \mathbf{z}_i\}) - g(\{\mathbf{z}_1, \ldots, \mathbf{z}_{i-1}\}).$$

By Lemma 3, $\delta_{i-1} \geq \delta_i$.

Because the point that was observed $\mathbf{r}_i$ that was closest to $\mathbf{z}_i^*$ was not added to S, we have the following inequalities

$$|\hat{g}(S \cup \{\mathbf{r}_i\}) - g(S \cup \{\mathbf{r}_i\})| \leq \epsilon_f$$
$$\hat{g}(S \cup \{\mathbf{r}_i\}) - \hat{g}(S \cup \{\mathbf{z}_b\}) \leq \epsilon_t$$
$$|g(S \cup \{\mathbf{r}_i\}) - g(S \cup \{\mathbf{z}_i^*\})| \leq \epsilon_r$$

where the last inequality is true for all $\mathbf{z}_i^*$ with probability $1 - \delta$. Using these inequalities, as shown more explicitly in the proof in the appendix, we get

$$g(S \cup \{\mathbf{z}_i^*\}) - g(S) \leq \delta_b + \epsilon_r + 2\epsilon_f + \epsilon_t$$

Algorithm 2 BlockGreedy: OnlineGreedy for Prototype Selection using a Block-Diagonal Approximation

r = block-size, with set of blocks $\mathcal{B}$, $S_0 \leftarrow \emptyset$
c, l book-keeping maps, with $(c(B), l(B)) = (\mathbf{z}, l)$ for $\mathbf{z}$ leading to smallest utility loss l if removed from block B.
$g_e \leftarrow 0$ is the incremental estimate of log-determinant
    for $t = 1 : b$ **do**, $S_t \leftarrow S_{t-1} \cup \{\mathbf{x}_t\}$
    while interacting, $t = b + 1, \ldots$ **do**
        if added b new prototypes since last clustering **then**
            cluster S_t into $\lfloor b/r \rfloor$ blocks with k-means,
            initialize with previous clustering; update c, l, g_e
        BlockGreedy-Swap($\mathbf{x}_t$)

Algorithm 3 BlockGreedy-Swap($\mathbf{x}$)

$B_1 \leftarrow$ get-block($\mathbf{x}$)   ▷ returns the nearest block to $\mathbf{x}$
$(\mathbf{z}_1, g_1) \leftarrow \underset{\mathbf{z} \in B_1}{\operatorname{argmax}} \, g(B_1 \backslash \{\mathbf{z}\} \cup \{\mathbf{x}\}) - g(B_1)$
if $g_e \neq 0$ and $\frac{g_1 - g_e}{g_e} < \epsilon_t$ **then**
    return with no update if low percentage improvement
$(B_2, g_2) \leftarrow \underset{B \in \mathcal{B} \backslash B_1}{\operatorname{argmax}} \, g(B_1 \cup \{\mathbf{x}\}) - l(B)$
if $g_1 < g_2$ **then**   ▷ remove point from same block
    $B_1 \leftarrow B_1 \backslash \{\mathbf{z}_1\} \cup \{\mathbf{x}\}$
    update $c(B_1)$ ▷ using Appendix E.3
    $g_e \leftarrow g_e + g_1$
else   ▷ remove point from a different block
    $B_1 \leftarrow B_1 \cup \{\mathbf{x}\}$
    $B_2 \leftarrow B_2 \setminus \{c(B_2)\}$
    update $c(B_1), c(B_2)$ ▷ using Appendix E.3
    $g_e \leftarrow g_e + g_2$

By the definition of submodularity, $g(S \cup S^*) - g(S) \leq \sum_{i=1}^{b} g(S \cup \{\mathbf{z}_i^*\}) - g(S))$.

Putting this all together, with probability $1 - \delta$,

$$g(S^*) \leq g(S \cup S^*)$$
$$\leq g(S) + \sum_{i=1}^{b} g(S \cup \{\mathbf{z}_i^*\}) - g(S))$$
$$\leq g(S) + \sum_{i=1}^{b} (\delta_b + \epsilon_r + 2\epsilon_f + \epsilon_t)$$
$$\leq g(S) + \left(\sum_{i=1}^{b} \delta_i \right) + b(\epsilon_r + 2\epsilon_f + \epsilon_t)$$
$$= g(S) + \sum_{i=1}^{b} g(\{s_1, \ldots, \mathbf{z}_i\}) - g(\{\mathbf{z}_1, \ldots, \mathbf{z}_{i-1}\})$$
$$+ b(\epsilon_r + 2\epsilon_f + \epsilon_t)$$
$$\leq g(S) + g(\{\mathbf{z}_1, \ldots, \mathbf{z}_b\}) + b(\epsilon_r + 2\epsilon_f + \epsilon_t)$$
$$\leq 2g(S_t) + b(\epsilon_r + 2\epsilon_f + \epsilon_t)$$

where the last inequality uses $g(S) \leq g(S_t)$ which follows from monotonicity. ∎

5. Block-diagonal approximation for efficient computation of the submodular function

The computation of the submodular function g is the critical bottleneck in OnlineGreedy and other incremental submodular maximization techniques. In this section, we propose a time and memory efficient greedy approach to computing a submodular function on $\mathbf{K}_S$, enabling each step of OnlineGreedy to cost $O(db)$, where d is the feature dimension and b is the budget size. The key insight is to take advantage of the block-diagonal structure of the kernel matrix, particularly due to the fact that the greedy algorithm intentionally selects diverse prototypes. Consequently, we can approximately cluster prototypes into small groups of size r, and perform updates on only these blocks.

Approximations to the kernel matrix have been extensively explored, but towards the aim of highly accurate approximations for use within prediction. These methods include low-rank approximations (Bach & Jordan, 2005), Nystrom methods (Drineas & Mahoney, 2005; Gittens & Mahoney, 2013) and a block-diagonal method for dense kernel matrices, focused on storage efficiency (Si et al., 2014). Because these approximations are used for prediction and because they are designed for a fixed batch of data and so do not take advantage of incrementally updating values, they are not sufficiently efficient for use on each step, and require at least $O(b^2)$ computation. For OnlineGreedy, however, we only need a more coarse approximation to $\mathbf{K}_S$ to enable effective prototype selection. By taking advantage of this fact, saving computation with incremental updating and using the fact that our kernel matrix is not dense—making it likely that many off-diagonal elements are near zero— we can reduce storage and computation to linear in b.

The key steps in the algorithm are to maintain a clustering of prototypes, compute all pairwise swaps between prototypes within a block—which is much more efficient than pairwise swaps between all prototypes— and finally perform a single swap between two blocks. The computational complexity of Algorithm 2 on each step is $O(bd + r^3)$ for block size r (see Appendix F for an in-depth explanation). We assume that, with a block-diagonal $\mathbf{K}_S$ with blocks $\mathcal{B}$, the submodular function separates into $g(S) = \sum_{B \in \mathcal{B}} g(B)$. For both the log-determinant and the trace of the inverse of $\mathbf{K}_S$, this is the case because the inverse of a block-diagonal matrix corresponds to a block-diagonal matrix of the inverses of these blocks. Therefore, $\log \det(\mathbf{K}_S) = \sum_{B \in \mathcal{B}} \log \det(\mathbf{K}_B)$.

We use this property to avoid all pairwise comparisons. If $\mathbf{x}$ is added to S, it gets clustered into its nearest block, based

(a) log det of **K**

(b) Estimate Accuracy, with $b = 200$

(c) Runtime with increasing b

Figure 1. Performance of BlockGreedy in Telemonitoring. Figure **(a)** shows the true log determinant of **K** for the prototypes selected by each algorithm. Our algorithm, BlockGreedy, achieves almost the same performance as FullGreedy, which uses no approximation to **K** to compute the log-determinant. Figure **(b)** shows the accuracy of the log determinant estimate as the block size increases. We can see that clustering is key, and that for smaller block sizes, comparing between blocks is key. Figure **(c)** shows the runtime of the main prototype selection competitors, FullGreedy and SieveStreaming, versus BlockGreedy with block size $r = 10$.

on distance to the mean of that cluster. To compute the log-determinant for the new S, we simply need to recompute the log-determinant for the modified block, as the log-determinant for the remaining blocks is unchanged. Therefore, if $\mathbf{K}_S$ really is block-diagonal, computing all pairwise swaps with $\mathbf{x}$ is equivalent to first computing the least useful point $\mathbf{z}$ in the closest cluster to $\mathbf{x}$, and then determining if $g(S \cup \{\mathbf{x}\})$ would be least reduced by removing $\mathbf{z}$ or removing the least useful prototype from another cluster. With some book-keeping, we maintain the least-useful prototype for each cluster, to avoid recomputing it each step.

6. Experiments

We empirically illustrate the accuracy and efficiency of our proposed method as compared to OnlineGreedy with no approximation to the submodular function (which we call Full Greedy), Sieve Streaming, and various naive versions of our algorithm. We also show this method can achieve reasonable regression accuracy as compared with KRLS (Engel et al., 2004). For these experiments we use four well known datasets: Boston Housing (Lichman, 2015), Parkinson's Telemonitoring (Tsanas et al., 2010), Sante Fe A (Weigend, 1994) and Census 1990 (Lichman, 2015). Further details about each dataset are in Appendix C. We use a Gaussian kernel for the first three datasets, and a Hamming distance kernel for Census, which has categorical features. To investigate the effect of the block-diagonal approximation, we select the log-determinant as the criterion, which is an instance of our criterion, and set $\lambda = 1$.

Quality of the log-determinant approximation.

We first investigate the quality of prototypes selection and their runtimes, depicted in Figure 1. We compare our al-

gorithm with the FullGreedy, SieveStreaming and a random prototype selection baseline. We also use variants of our algorithm including without clustering—naively dividing prototypes into equal-sized blocks—and one where we only consider replacement in the closest block. We include these variants to indicate the importance of clustering and of searching between blocks as well within blocks, in BlockGreedy. For all experiments on maximization quality, we use percentage gain with a threshold of $\epsilon_t = 0.001$.

We plot the log determinant with increasing samples, in Figure 1(a). Experiments on the other datasets are included in Appendix C. BlockGreedy maximizes the submodular function within 1% of the FullGreedy method. Though BlockGreedy achieves nearly as high a log determinant value, we can see that its approximation of the log determinant is an overestimate for this small block size, $r = 5$.

Our next experiment, therefore, focuses on the estimate accuracy of BlockGreedy with increasing block size, in Figure 1(b). The accuracy is computed by $1 - \left| \frac{g_{actual} - g_{estimate}}{g_{actual}} \right|$. We can see our algorithm, BlockGreedy performs much better as compared to the other variants, ranging in accuracy from 0.82 to 0.99. This suggests that one can choose reasonably small block sizes, without incurring a significant penalty in maximizing the log determinant. In Figure 1(a), the estimate is inaccurate by about 20%, but follows the same trend of the full log determinant and picks similar prototypes to those chosen by FullGreedy.

The runtime of our algorithm should be much less than that of FullGreedy, and memory overhead much less than SieveStreaming. In Figure 1(c), we can see our method scales much better than FullGreedy and even has gains in speed over SieveStreaming. Though not shown, the number of sieves generated by SieveStreaming is large, in many instances well over 600, introducing a significant amount of

(a) Boston housing   (b) Telemonitoring   (c) Santa Fe Data Set A

Figure 2. Figure **(a)** is the learning curve on Bostondata, averaged over 50 runs, with $\eta = 0.01$. On average, KRLS uses 116.46 prototypes. Figure **(b)** is the learning curve on the Telemonitoring data set, over 50 runs, with $b = 500$ and $\eta = 0.001$. Figure **(c)** plots the predicted values of our algorithm and true values. The regularizer $\eta = 0.001$, and the utility threshold is $\epsilon_t = 0.0001$.

overhead. Overall, by taking advantage of the block structure of the kernel matrix, our algorithm obtains significant runtime and memory improvements, while also producing a highly accurate estimate of the log determinant.

Learning performance for regression problems.

While the maximization of the submodular function is useful in creating a diverse collection of prototypes, ultimately we would like to use these representations for prediction. In Figure 2, we show the effectiveness of solving $(\mathbf{K}_S + \eta \mathbf{I})\mathbf{w} = \mathbf{y}$ for the three regression datasets, by using our algorithm to select prototypes for $\mathbf{K}_S$. For all regression experiments, we use a threshold of $\epsilon_t = 0.01$ unless otherwise specified.

For Boston housing data, in figure 2(a), we see that Block-Greedy can perform almost as well as FullGreedy and SieveStreaming, and outperforms KRLS at early learning and finally converges to almost same performance. We set the parameters for KRLS using the same parameter settings for this dataset as in their paper (Engel et al., 2004). For our algorithms we set the budget size to $b = 80$, which is smaller than what KRLS used, and chose a block size of 4. We also have lower learning variance than KRLS, likely because we use explicit regularization, whereas KRLS uses its prototype selection mechanism for regularization.

On the Telemonitoring dataset, the competitive algorithms all perform equally well, reaching a RMSE of approximately 4.797. BlockGreedy, however, uses significantly less computation for selecting prototypes. We used a budget of $b = 500$, and a block size of $r = 25$; a block size of $r = 5$ for this many prototypes impacted the log determinant estimation enough that it was only able to reach a RMSE of about 5.2. With the larger block size, Block-Greedy obtained a log determinant value within 0.5% of FullGreedy.

On the benchmark time series data set *Santa Fe Data Set A*, we train on the first 1000 time steps in the series and predict the next 100 steps, calculating the normalized MSE (NMSE), as stated in the original competition. We set the width parameter and budget size to that used with KRLS after one iteration on the training set. The NMSE of our algorithm and KRLS were 0.0434 and 0.026 respectively. While our method performs worse, note that KRLS actually runs on 6×1000 samples according to its description (Engel et al., 2004), but with 1000 samples it performs worse with a NMSE of 0.0661. We demonstrate the 100-step forecast with BlockGreedy, in Figure 2(c); we include forecast plots for the other algorithms in Figure 4, Appendix C.4.

7. Conclusion

We developed a memory and computation efficient incremental algorithm, called BlockGreedy, to select centers for kernel representations in a continual learning setting. We derived a criterion for prototype selection, and showed that the log-determinant is an instance of this criterion. We extended results from streaming submodular maximization, to obtain an approximation ratio for OnlineGreedy. We then derived the efficient variant, BlockGreedy, to take advantage of the block-diagonal structure of the kernel matrix, which enables separability of the criteria and faster local computations. We demonstrated that, by taking advantage of this structure, BlockGreedy can significantly reduce computation without incurring much penalty in maximizing the log-determinant and maintaining competitive prediction performance. Our goal within continual learning was to provide a principled, near-linear time algorithm for prototype selection, in terms of the number of prototypes. We believe that BlockGreedy provides one of the first such algorithms, and is an important step towards effective kernel representations for continual learning settings, like online learning and reinforcement learning.

References

Argyriou, A. and Dinuzzo, F. A Unifying View of Representer Theorems. In *International Conference on Machine Learning*, 2014.

Bach, F. R. and Jordan, M. I. Predictive low-rank decomposition for kernel methods. In *International Conference on Machine Learning*, 2005.

Badanidiyuru, A., Mirzasoleiman, B., Karbasi, A., and Krause, A. Streaming submodular maximization: massive data summarization on the fly. *Conference on Knowledge Discovery and Data Mining*, 2014.

Brookes, M. *Matrix reference manual*. Imperial College London, 2004.

Chen, B., Zhao, S., Zhu, P., and Principe, J. C. Quantized Kernel Recursive Least Squares Algorithm. *IEEE Transactions on Neural Networks and Learning Systems*, 2013.

Cheng, C.-A. and Boots, B. Incremental Variational Sparse Gaussian Process Regression. *Advances in Neural Information Processing Systems*, 2016.

Cheng, L., Vishwanathan, S. V. N., Schuurmans, D., Wang, S., and Caelli, S. W. Implicit Online Learning with Kernels. In *Advances in Neural Information Processing Systems*, 2007.

Csató, L. and Opper, M. Sparse On-Line Gaussian Processes. *dx.doi.org*, 2006.

Dai, B., Xie, B., He, N., Liang, Y., Raj, A., Balcan, M.-F. F., and Song, L. Scalable Kernel Methods via Doubly Stochastic Gradients. *Advances in Neural Information Processing Systems*, 2014.

Das, A. and Kempe, D. Submodular meets Spectral: Greedy Algorithms for Subset Selection, Sparse Approximation and Dictionary Selection. *arXiv.org*, 2011.

Drineas, P. and Mahoney, M. W. On the Nyström Method for Approximating a Gram Matrix for Improved Kernel-Based Learning. *Journal of Machine Learning Research*, 2005.

Engel, Y., Mannor, S., and Meir, R. The kernel recursive least-squares algorithm. *IEEE Transactions on Signal Processing*, 2004.

Gittens, A. and Mahoney, M. W. Revisiting the Nystrom method for improved large-scale machine learning. In *International Conference on Machine Learning*, 2013.

Guha, S., Meyerson, A., Mishra, N., Motwani, R., and O'Callaghan, L. Clustering Data Streams: Theory and Practice. *IEEE Transaction on Knowledge and Data Engineering*, 2003.

Huang, P. S., Avron, H., and Sainath, T. N. Kernel methods match deep neural networks on timit. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014.

Kivinen, J., Smola, A., and Williamson, R. C. Online learning with kernels. *IEEE Transactions on Signal Processing*, 2010.

Krause, A. and Gomes, R. G. Budgeted nonparametric learning from data streams. In *International Conference on Machine Learning*, 2010.

Krause, A., McMahon, H. B., Guestrin, C., and Gupta, A. Robust Submodular Observation Selection. *Journal of Machine Learning Research*, 2008a.

Krause, A., Singh, A. P., and Guestrin, C. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *Journal of Machine Learning Research*, 2008b.

Li, C., Jegelka, S., and Sra, S. Fast DPP Sampling for Nystrom with Application to Kernel Methods. In *International Conference on Machine Learning*, 2016.

Lichman, M. *UCI machine learning repository*. URL http://archive. ics. uci. edu/ml, 2015.

Lu, Z., May, A., Liu, K., Garakani, A. B., Guo, D., Bellet, A., Fan, L., Collins, M., Kingsbury, B., Picheny, M., and Sha, F. How to Scale Up Kernel Methods to Be As Good As Deep Neural Nets. *CoRR abs/1202.6504*, 2014.

Mairal, J., Koniusz, P., Harchaoui, Z., and Schmid, C. Convolutional Kernel Networks. *NIPS*, 2014.

Matic, I. Inequalities with determinants of perturbed positive matrices. *Linear Algebra and its Applications*, 2014.

Park, J. and Sandberg, I. Universal Approximation Using Radial-Basis-Function Networks. *Neural Computation*, 1991.

Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2007.

Richard, C., Bermudez, J. C. M., and Honeine, P. Online Prediction of Time Series Data With Kernels. *IEEE Transactions on Signal Processing*, 2009.

Rudi, A., Camoriano, R., and Rosasco, L. Less is More: Nyström Computational Regularization. *Advances in Neural Information Processing Systems*, 2015.

Schraudolph, N. N., Smola, A. J., and Joachims, T. Step size adaptation in reproducing kernel Hilbert space. *Journal of Machine Learning Research*, 2006.

Seeger, M. Greedy Forward Selection in the Informative Vector Machine. 2004.

Seeger, M., Williams, C., and Lawrence, N. Fast Forward Selection to Speed Up Sparse Gaussian Process Regression. *Artificial Intelligence and Statistics*, 2003.

Si, S., Hsieh, C.-J., and Dhillon, I. Memory efficient kernel approximation. In *International Conference on Machine Learning*, 2014.

Tsanas, A., Little, M. A., and McSharry, P. E. Accurate Telemonitoring of Parkinson's Disease Progression by Noninvasive Speech Tests. *IEEE transactions on Biomedical Engineering*, 2010.

Van Vaerenbergh, S. and Santamaria, I. A comparative study of kernel adaptive filtering algorithms. In *Digital Signal Processing and Signal Processing Education Meeting*, 2013.

Van Vaerenbergh, S., Santamaría, I., Liu, W., and Principe, J. C. Fixed-budget kernel recursive least-squares. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2010.

Weigend, A. S. Time Series Prediction: Forecasting the Future and Understanding the Past. *Santa Fe Institute Studies in the Sciences of Complexity*, 1994.