

Prediction and Discovery of Users' Desktop Behavior

Omid Madani and Hung Bui and Eric Yeh

Artificial Intelligence Center, SRI International
333 Ravenswood Ave., Menlo Park, CA 94025

Abstract

We investigate prediction and discovery of user desktop activities. The techniques we explore are unsupervised. In the first part of the paper, we show that efficient many-class learning can perform well for action prediction in the Unix domain, significantly improving over previously published results. This finding is promising for various human-computer interaction scenarios where rich predictive features of different types may be available and where there can be substantial nonstationarity. In the second part, we briefly explore techniques for extracting salient activity patterns or *motifs*. Such motifs are useful in obtaining insights into user behavior, automated discovery of (often interleaved) high-level tasks, and activity tracking and prediction.

1 Introduction

An exciting and promising domain for machine learning continues to be the area of action monitoring and personalization. Adaptive systems find applications in task completion, for instance in aiding users' desktop activity, or in reminding users of actions they may have forgotten (assisted living) or proposing alternative possibilities. In this paper, we investigate user action prediction as well as discovery of salient activity. Our approaches are unsupervised, in that the user does not explicitly try to teach the system about her activity. The system simply observes and learns to predict actions or discovers salient patterns (motifs). As in (Davison and Hirsh 1998; Korvemaker and Greiner 2000), our experiments will be primarily in the Unix domain (Greenberg 1988), as much data on a variety of users is available, and we can compare prediction accuracy. We also report on logs of our own desktop activity using the TaskTracer system (Dragunov et al. 2005). The paper is divided into two parts.

1.1 Action Prediction. Here, we focus on the problem of predicting the entire command line that the user would want to type next. Depending on the attributes of the context, such as time of day, current working directory, and recently performed actions, the system may predict that the next command will be “make”, or “latex paper.tex”, or “cd courses”, and so on. The user interface can depend on the particularities of the task. For instance, as explained in (Korvemaker

and Greiner 2000), the top five predictions of the system can be tied to the function keys F1 through F5. If the user observes the correct command suggested (e.g., on the top bar of the window), a simple key press executes the action. This can nicely complement other Unix facilities that aid in typing commands. We note that while our experiments are in the Unix domain, similar problems arise in other desktop interaction contexts. For instance, in the Windows domain, the problem can be predicting the next directory to which or from which the user will choose to save an email attachment or load a file (Bao, Herlocker, and Dieterich 2006). Other domains include devices with a limited interface, such as cell phones.

1.1.1 Challenges. The prediction task entails a number of challenges, including (1) high dimensionality and in particular many classes, (2) space and time efficiency, and (3) nonstationarity. We seek algorithms that can capture context well. This means the effective aggregation of the predictions of a rich set of features (predictors). A core aspect that distinguishes our approach is viewing the task as a *many-class* learning problem (multiclass learning with many classes: 100s, 1000s, . . .). The different number of items to predict, entire commands or parameters and so on, ranges in hundreds in our experiments. We employ recently developed efficient *indexing* algorithms (Madani and Connor 2008; Madani and Huang 2008). Efficiency is paramount in this domain: the system must quickly respond and remain adaptive. The algorithms we describe are efficient both in space consumption and in time. As we will see, the prediction problem is significantly nonstationary. In this paper, we evaluate the many-class algorithms in this setting for the first time (as opposed to the more common batch setting). Due to nonstationarity, an important question is whether a learner has sufficient time (learning period) to be able to learn good aggregation of the many features. Indeed (Korvemaker and Greiner 2000), after trying a number of context attributes to improve prediction over using a basic method (as we explain), and failing to improve accuracy, conclude that all the prediction signal may have been gleaned, yet we show that via using improved learning techniques we can gain substantially (an average of 4% to 5% in absolute accuracy improvement, or about 10% relative, over 168 users). We also find that the one-versus-rest linear SVM has very poor accuracy in this domain.

1.2 Motif Discovery. In the second part of the paper, we briefly present our efforts on extracting *salient* repeated patterns, in the form of directed subgraphs, from the action logs (both Unix and desktop). Such patterns provide insights into user behavior, can be used as features for prediction, and can be a basis for learning of higher-level user tasks.

Paper Organization. The next section describes the problem domain, choice of features, algorithms, and evaluation methods. Section 3 presents a variety of experiments and comparisons, Section 4 presents our motif discovery work, Section 5 discusses related work, and Section 6 concludes with future work. An expanded version of this paper with further experiments is in preparation (Madani, Bui, and Yeh 2009).

2 Preliminaries

Our setting is standard multiclass supervised learning, but with many classes and nonstationarity. A learning problem consists of a sequence of instances, each training instance specified by a vector of feature values, x , and the class that the instance belongs to y_x (the *positive class*). We use x to refer to the instance itself as well. Given an instance, a *negative class* is any class $c \neq y_x$. x_f denotes the value of feature f in the vector x . We enforce that $x_f \geq 0$. If $x_f > 0$, we say feature f is *active* in instance x , and denote this aspect by $f \in x$. The number of active features in x is denoted by $|x|$.

2.1 Data sets and Tasks

The bulk of our experiments is performed on a data set collected by Greenberg on Unix usage (Greenberg 1988). This data set is fairly large, collected on 168 users over 2 to 6 months. There are four user types: 52 computer scientists, 36 expert programmers, 55 novice programmers, and 25 non-programmers. This data set also allows us to compare to previous published results.

We use the terminology of (Korvemaker and Greiner 2000): a (full) command is the entirety of what is entered, this includes the “stub”, meaning the “executable” (or action) part, and possibly options and parameters (file and directory names). Thus, in “ls -l home”, “ls” is the stub part, “home” is the parameter, and the command is “ls -l home”. We will focus on the task of learning to predict the (full) commands, as in (Korvemaker and Greiner 2000). For this task, the number of unique classes (commands) on average per user is roughly 470. As one may expect, this average is highest for computer scientists as a group (around 700 on average), next is experienced programmers (500), and novice programmers and non-programmers have about the same (300). It was found that computer scientists were the hardest to predict as a group (Korvemaker and Greiner 2000). As in (Korvemaker and Greiner 2000), over all the users, we obtain 303,628 episodes (commands entered).

2.2 The Choice of Features and Representation

We experimented with the following feature types, which we break into two broad categories. *Action features* reflect what the user has done recently. The ones we used are the

commands typed at times $t - 1$ and $t - 2$,¹ as well as only the stub and only parameter portions of the command, at time $t - 1$. We also found the *start-session* feature to be useful (each user’s log is broken into many sessions in which the user begins the session, and after some interaction, exits the session). The start-session feature was treated like other action features. Thus, after starting, the start session feature would be the “command” taken at time $t - 1$ and after one command entered, it would be the command taken at time $t - 2$.

The other type of features may be called *State Features*, i.e., those that reflect the “state” the user or the system is in. State features do not change as quickly as action features. We used the current working directory as one state feature. Importantly, we also used a “default” or an *always-active-feature*, a feature with value of 1 in every episode and that would be updated in every update (but not necessarily in every episode). This feature has an effect similar to the LIFO strategy (see Section 2.4). Episodes have less than 10 features active on average. The very first episode of a user has three features active: the always-active feature, start session feature, and a feature indicating that the last command had no parameter (the “NULL” parameter).²

We did not attempt to predict the start of session nor the exiting action. The recorded logs also indicated whether an error occurred. A significant portion of the commands led to errors (e.g., mistyped commands) (about 5% macro averaged over users). We did not treat them differently. These decisions allowed to us to compare to the results of (Korvemaker and Greiner 2000).

2.3 Online Evaluation

All the algorithms we evaluate output a ranking of their predictions. As in (Korvemaker and Greiner 2000; Davison and Hirsh 1998), unless specified otherwise, we report on the *cumulative online* accuracy (ranking) performance of R_1 (standard accuracy, or one minus zero-one error) and R_5 (accuracy in top five predictions), computed for each user, then averaged over all the 168 users (macro averaged). For this evaluation, for each user, the sequence of episodes (instances or commands) is ordered naturally in the order they were typed. Formally, let k_{x_i} be the rank of the positive class y_{x_i} for the i -th instance x_i in the sequence. Let $\mathbb{I}\{k_{x_i} \leq k\} = 1$ iff $k_{x_i} \leq k$, and 0 otherwise (Iverson bracket). Then R_k (R_1 or R_5) for a given user with M instances is

$$R_k = \frac{1}{M} \sum_{1 \leq i \leq M} \mathbb{I}\{k_{x_i} \leq k\} \quad (1)$$

On each instance, first the system is evaluated (predicts using the features of the instance), then the system trains on that instance (the true class is revealed). The algorithms we present in Section 2.4 perform a simple efficient prediction

¹As separate features, i.e., even if the same command was typed again, it has a different feature id for times $t - 1$ and $t - 2$.

²We could also have added the home directory.

and possible update on each instance. We note that the system always fails on the first instance, and more generally on any instance for which the true class has not been seen before. As in (Korvemaker and Greiner 2000), such instances are included in the evaluation. On average per user, about 17% of commands are not seen before. This number goes to over 25% for parameter portion of commands, and down to 6% for stubs. We will also report on some variations, such as when the instances are permuted, to obtain insights on algorithms’ performances in the more common “stationary” evaluation setting.

2.4 Algorithms

The main learning algorithm that we propose for the prediction task, shown Figure 1, employs exponential moving average updating, and we refer to it as EMA (“Emma”). On every instance (episode), the algorithm first predicts the class (in our case, the user’s next command line) and, if margin threshold is not met, updates each active feature using exponential moving average updating. The connection (prediction) weight from feature f to class c is denoted by $w_{f,c}$. Updates are kept efficient as each active feature resets weights (connections) that fall below a threshold w_{min} . In our experiments, w_{min} is set to 0.01; thus, the maximum out-degree of a feature, denoted d , is 100. The connections of a feature are implemented via a dynamic sorted linked list. The first time a feature is seen (in some episode), it is not connected to any class (all its connection weights are implicitly 0). We have termed the learned representation an *index*, i.e., a mapping that connects each feature to a relatively small subset of the classes (the features “index” the classes). Both prediction and updating on an instance x take time $O(d|x| \log(d|x|))$. Several properties of EMA were explored in (Madani and Huang 2008). It was shown that EMA updating is equivalent to a quadratic loss minimization for each feature, and a formulation for numeric feature values, as given in Figure 1, was derived. An update is performed only if a margin threshold δ_m is not met (a kind of mistake-driven updating). This leads to down weighing the votes of redundant features, more effective aggregation, and ultimately better generalization (Madani and Connor 2008).

We compare EMA against the method used in (Korvemaker and Greiner 2000) as well as linear SVMs, and another indexing variant OZ (Madani and Huang 2008). In the approach of (Korvemaker and Greiner 2000), a restricted version of EMA updating was deployed, which was referred to as the alpha updating rule at the time, or AUR. AUR can also be viewed as a nonstationary version of the bigram method in statistical language modeling. A similar strategy was used in (Davison and Hirsh 1998), but for the task of stub prediction. In AUR, only the last command is used as a predictor, with one exception: if that command does not give at least five predictions, a default predictor is used to fill in the remaining of the five slots. Whenever a command appears as a feature, it is updated, and the default predictor is updated in every episode (using exponential moving average). Thus, the differences with our presentation of EMA are that we use multiple features and aggregate their votes (their method did not sum, only merge the predictions if

EMA($x, y_x, \beta, \delta_m, w_{min}$)

1. $\forall c, s_c \leftarrow \sum_{f \in x} x_f w_{f,c}$ /* Score classes for ranking/prediction */
2. $\delta_x \leftarrow s_{y_x} - s_{c'}$, where $s_{c'} \leftarrow \max_{c \neq y_x} s_c$ /* compute margin */
3. if ($\delta_x < \delta_m$), then $\forall f \in x$ do: /* If margin not met, update */
 - /* All active features’ connections are decayed, then the connections to true class is boosted */
 - 3.1 $\forall c, w_{f,c} \leftarrow (1 - x_f^2 \beta) w_{f,c}$ /* $0 < x_f \leq 1$ */
 - 3.2 $w_{f,y_x} \leftarrow w_{f,y_x} + x_f \beta$ /* Boost connection to true class */
 - 3.3 If $w_{f,c} < w_{min}$, then /* drop small weights */
 - $w_{f,c} \leftarrow 0$

Figure 1: The EMA (“Emma”) learning algorithm, which uses *exponential moving average* updating and a margin threshold. β is a learning rate or a “boost” amount, $0 < \beta \leq 1$, and x_f is the “activity” level of active feature f , $0 < x_f \leq 1$. The end effect after an update is that the weight of the connection of feature f to the positive class, w_{f,y_x} , is strengthened. Other connections are weakened and possibly zeroed (dropped).

	$R1$	$R5$	$R1_{>1k}$	$R5_{>1k}$
LIFO	0.075 \pm 0.05	0.42 \pm 0.15	0.07 \pm 0.04	0.40 \pm 0.15
AUR	0.28 \pm 0.12	0.47 \pm 0.14	0.28 \pm 0.12	0.47 \pm 0.14
EMA	0.30 \pm 0.11	0.51 \pm 0.13	0.30 \pm 0.11	0.52 \pm 0.13

Table 1: Accuracies on full command prediction. $R1_{>1k}$ and $R5_{>1k}$ are averages on users with no less than 1000 episodes (EMA: $\beta = 0.15, \delta_m = 0.15, w_{min} = 0.01$).

need be), we update in a mistake-driven manner (in particular we use a margin threshold), we drop weak edges, and we l_2 normalize the feature vectors. We also compare against a last-in-first-out strategy, or LIFO. LIFO keeps track of the last five *unique* commands and reports them in that (reverse chronological) order, so that the command typed last (time $t - 1$) is reported first.³

3 Experiments

Table 1 shows the performance comparisons between LIFO, AUR, and EMA. For all these methods, an entire evaluation on all 168 users takes less than 2 minutes on a laptop. We observe that the effective aggregation of predictors (or capturing more context) can lead to a substantial boost in accuracy, in particular in $R5$. The performance on users with more than 1000 instances appears to lead to some improvement for EMA (but not for AUR)⁴, over those users with fewer than 1000. We note that Korvemaker and Greiner tried a number of ways and features to improve on AUR, but their methods did not lead to a performance gain (Korvemaker and Greiner 2000). In Figure 2, the performance on each

³Another similar baseline is reporting the five most frequent commands seen so far. As (Korvemaker and Greiner 2000) show, that strategy performs substantially worse than LIFO (several percentage points below in accuracy), underscoring the nonstationarity aspect.

⁴For AUR, we obtained the same accuracies of (Korvemaker and Greiner 2000).




Figure 2: The spread of performance over users. For each user, the x-axis is the performance of AUR (R1 or R5), the y-axis is the difference from EMA (above 0 means higher performance for EMA).

user is depicted by a point where the x-coordinate is R1 (or R5) using AUR, and the y-coordinate is that same value subtracted from R1 (or R5) obtained using EMA. We observe that R_5 values in particular are significantly improved using EMA, and the improvements tend to be higher (in absolute as well as relative value) with lower absolute value of the performance. The values for users with fewer than 1000 instances shows somewhat higher spread, as may be expected. We compared the number of wins, when results on the same user are paired, and performed a sign test. On R1, EMA wins over AUR on 141 of the users, loses on 26 users, and ties on 1 (Figure 2). On R5, winning is more robust: EMA wins in 162 cases and loses in 6. Both comparisons are significant with over 99.9% confidence (P value is < 0.001).

3.1 Ablation Experiments on EMA

Figure 3 shows (macro) average R5 performance as a function of learning rate and margin. We notice the performances are fairly close: the algorithm is not heavily dependent on the parameters. It is also interesting to note that relatively high learning rates of 0.1 and above give the best or very




Figure 3: Plots of R5 performance (macro average over users) as a function of choices for margin threshold and the learning rate.

good results here. In previous studies in text categorization (Madani and Huang 2008), lower learning rates (of 0.05 or 0.01 and below) gave the best results. When we compare the best overall R5 average for learning rate of 0.15 versus 0.05, we obtain 120 wins (for learning rate of 0.15), 41 losses and 7 ties (where the average is respectively R5 of 0.51 for learning rate of 0.15 and 0.50 for learning rate of 0.05). As might be expected, the best results are obtained when the margin (threshold) is not at the extremes of 0 (pure mistake-driven “lazy” updating) or very high (always update). If we include more features, such as stub and parameters from time $t - 2$ or features from earlier time points, tending to increase redundancy and uninformative, performance somewhat degrades (the average remains 0.5 or around it), and the selection of margin becomes more important. It may be possible to adjust (learn) the learning rate or margin over time as a function of user behavior for improved performance.

With the default parameters of 0.15 for both learning rate and margin, we raised the minimum weight threshold w_{min} to 0.05 and 0.1 (from default of 0.01), and respectively obtained R_5 of 0.509 (small degradation) and 0.45 (substantial degradation).

3.2 Feature Utilities

In the results given here, the default parameters are used and all features are available (as explained in Section 2.2) except for those that we explicitly say we remove. The performance is fairly robust to removal of various feature types: the remaining feature types tend to compensate. All the features tend to help the average performance somewhat. Removing the stub or the (full) command at time $t - 1$ yields the largest drop in performance, leading to just over 0.50 average R_5 . If we remove both, we get an $R_5 = 0.486$. The other features in order of importance are current directory, always active, start session, and parameter at time $t - 1$. Removing any such type of feature results in degradation of about 0.005 (from the maximum of just over 0.51).




Figure 4: The need for sustained learning: Performance plots (R_5 for scientist 52) obtained from continued learning, LIFO, as well as cases in which (EMA) learning is stopped after a certain episode (1000, 1500, and 2000) (top: cumulative performance, bottom: moving average with β of 0.01). Continued learning is required to sustain prediction performance.

3.3 The Need for Continued Adaptation

Figure 4 shows R_5 performances as a function of time, averaged over the first t episodes, and the moving average, for scientist 52, for whom we have about 8000 total episodes. While the learning curve seems to reach local maxima at say around 1000 to 2000, we see the need for continued learning to *sustain* the performance: if we stop the learning, the performance curve takes a downward turn. The system’s performance eventually degrades to below that of LIFO if learning is stopped.

3.4 Comparisons with other Methods

Davison and Hirsch performed comparisons on the task of stub prediction (Davison and Hirsh 1998) and showed that their AUR rule for stubs (using the predictions of the previous stub) outperformed batch learning algorithms such as decision trees and naive Bayes as well as stationary variants of probability computation on their task. EMA outper-

random split	SVM	EMA	OOZ
R1	0.057	0.247	0.259
R5	0.306	0.474	0.484
chronological	SVM	EMA	OOZ
R1	0.01	0.210	0.187
R5	0.215	0.404	0.391

Figure 5: Comparisons between EMA, one-versus-rest linear SVM, and OOZ, on 21 selected users who had at least 1000 episodes. For all the algorithms, several parameters were tried (see text) and the best, on R_5 , is reported, with 80% for training, 20% test. (top) Traditional “batch” evaluation setting wherein instances are randomly permuted before splitting data. (bottom) Chronological split, where the first 80% is used for training.

forms AUR on stub prediction as well (about 3% improvement in R_5). Here we compare EMA with another online feature-focus variant OOZ (Madani and Huang 2008), as well as one-versus-rest (binary) linear SVM training (Figure 5). Over the 168 users, OOZ obtained a maximum accuracy of around 0.481 on R_5 , with the best parameter choice of margin close to 0, and β of 0.05 or 0.15, significantly underperforming EMA. This underscores the nonstationarity of the problem: while OOZ is an online method, the emphasis on hinge-loss minimization may not be the appropriate learning bias in this nonstationary task. EMA, with its weakening of all connections, may be more appropriate (unlike OOZ, which weakens certain most-violating connections): EMA can quickly forget associations that become old (e.g., do not repeat).

To further test this aspect, as well as to compare to one-versus-rest SVMs (substantially slower), we chose 21 users that had a number of episodes exceeding 1000 and compared performance in a *batch* setting, as well as the online or chronological setting. In the batch setting, each user’s episodes are randomly permuted, 20% are held out for evaluation, and the algorithms are trained on the remaining 80%. In the online setting, the first 80% of episodes are kept for training (in chronological order), the remaining for test. These instances are presented to the online learners, EMA and OOZ, in order. For SVM training the order does not make a difference. In all cases, as before, the feature vectors are l_2 normalized. We use the same feature representations (default setting) in all experiments. For all the algorithms, we evaluated under many different parameters to get their best performance. For the linear SVM, the regularization parameter was picked from the set $C \in \{0.5, 1, 5, 10, 20, 100\}$. For both OOZ and EMA, we tested learning rates in $\{0.01, 0.05, 0.1, 0.15, 0.25, 0.5\}$ and margin thresholds in $\{0.01, 0.05, 0.15, 0.25, 0.5\}$. We see that one-versus-rest SVM underperforms substantially, under both conditions. As explained in (Madani and Connor 2008), one-versus-rest can do a relatively poor job of ranking classes for a given instance, when there are many classes. SVM is also a batch algorithm. OOZ performs better than EMA under the batch setting (similar to findings of (Madani and Huang 2008)), but lags on the chronological splits.

4 Motif Discovery

In addition to stub and full command prediction, we attempt to induce *motifs*, meaningful sequences of related events that can correspond to coherent high-level tasks. This section is short, only describing and briefly demonstrating the approach. Our data for this experiment is taken from the selected 21 Unix activity logs (Section 3.4), and from two users’ Windows desktop activities, recorded over a period of three months using the TaskTracer system (Dragunov et al. 2005). In TaskTracer the events have a form similar to the Unix commands, consisting of event names (or stubs) such as “SendEmai” or “MSWord.Open” along with arguments such as the email’s recipient, and the file being opened. Our intent is to build up higher-level models of user activities from these sequences, and to eventually leverage these motifs for other purposes such as task tracking.

Central to the discovery of motifs is a measurement of association between pairs of events. We adapt an often-used measurement of word association in the NLP domain, namely, the pointwise mutual information between two words $\log \frac{\Pr(b|a)}{\Pr(b)}$ ((Manning and Schütze 1999), see also (Chambers and Jurafsky 2008)). This basic measurement needs to be extended for our domain so that (1) the order of the events matters, (2) events that are temporally close or share the same argument values are more likely to be associated.

We propose the following extension to the pointwise mutual information score. Let ω be a random window (of fixed size $d \geq 2$) of events X_i, \dots, X_{i+d-1} . Note that each event X_i is pair of stub and argument list (T_i, G_i) . We define the following random variables. $S(\omega)$ is the stub (event name) of the first event in ω . For each stub a , $C_a(\omega)$ is a random variable representing the cohesiveness of a with respect to the beginning event of ω :

$$C_a(\omega) = \sum_{k=1}^{d-1} \left(1 - \frac{k-1}{d-1}\right) \text{Overlap}(G_{i+k}, G_i) \mathbb{I}\{T_{i+k} = a\}$$

In the above, the first term penalizes events further away in the window, the second term, $\text{Overlap}(G, G') = \frac{|G \cap G'|}{|G \cup G'|}$, measures how similar the two argument lists are, and the third $\mathbb{I}\{T_{i+k} = a\}$ is the Iverson bracket (value of 1 if $T_{i+k} = a$, otherwise 0). Our measure of association for an ordered pair of stubs (a, b) can be taken as

$$\text{Assoc}(a, b) = \log \frac{\mathbb{E}[C_b(\omega) | S(\omega) = a]}{\mathbb{E}[C_b(\omega)]}$$

To gain further insight about this association measure, we first note that if the random variable C_a is boolean (i.e., indicator function), then the expectations reduce to probabilities and we obtain the familiar form for pointwise mutual information. Further, if a and b are uncorrelated, one might conjecture that knowing the window begins with a would not change the expected value of C_b , yielding an association score of zero. On the other hand, if b tends to follow a , knowing the same condition would increase the expected value of C_b , yielding a positive association score. In practice, the score is computed (estimated) by replacing the expectation with the empirical expectation evaluated from the

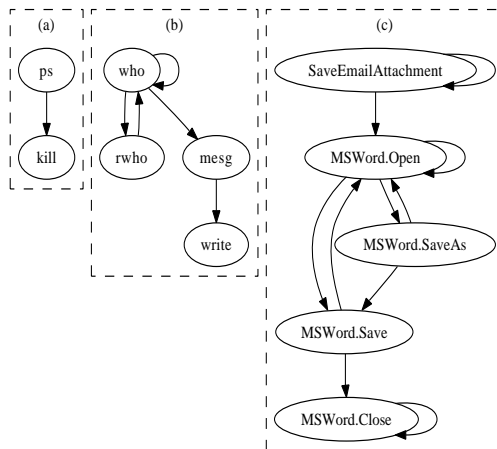


Figure 6: Example motifs drawn over a few users.

data. Finally, we note that the proposed association score is not symmetric as the order of the events is important.

Using the above, an association graph is constructed for each log, with nodes consisting of all unique stubs. A directed edge is drawn from stub a to b if $\text{Assoc}(a, b)$ exceeds a threshold. Motifs are extracted by finding subgraphs of the association graphs. The window size for the Unix dataset is set to 3 to reflect the shorter nature of command motifs, and is set to 50 for the Desktop data. Rarely occurring (fewer than five times) stubs are filtered out and not included in the graphs.

In general, the motifs matched to plausible user activity sequences. From sample motifs in Figure 6, we can identify several high-level user tasks, such as listing and killing processes in (a) or sending quick messages to other users (b) in the Unix domain. In the Desktop domain, (c) is an example of saving several MS Word attachments and operating over them. In comparison with Desktop motifs, those from the Unix set tended to be shorter, and in the case of the computer scientists, appeared more arbitrary. A possible explanation would be that the Unix data consisted of command line interactions, and not activity inside applications, such as Emacs or mail, whereas the Desktop activity record does contain events from analogous applications, such as Microsoft Word and Outlook. In addition, we find that argument sharing is an important feature that should be utilized by the association score. Dropping the Overlap term in the score leads to densely connected association graphs where individual motifs are no longer well separated.

5 Related Work

Modeling user activities on the desktop has been an active topic of research in the AI and user modeling community. It is generally accepted that good predictive models of user activities play a central part in building an intelligent adaptive interface that can potentially help to increase user productivity. The availability of the UNIX data (Greenberg 1988) has led to a number of efforts in building predictive models

of the data as well as facilitated direct and objective comparisons among different algorithms (Davison and Hirsh 1998; Korvemaker and Greiner 2000). These authors have also observed the nonstationarity nature of the task. For the related problem of automated email foldering, see also (Segal and Kephart 2000) on the importance of incremental learning, and (Bekkerman and McCallum 2004) on the importance of taking account of the nonstationarities in evaluation. Work on modeling user interactions with Windows can be traced to the LookOut system (Horvitz 1999), which can observe user email and calendar activities and attempt to learn a model of calendar-related emails. More recent work in this area includes the TaskTracer system (Dragunov et al. 2005) and BusyBody system (Kapoor and Horvitz 2007). These systems can capture a wide range of Windows and application events. The recorded events have been used for training various prediction tasks such as folder prediction (Bao, Herlocker, and Dietterich 2006), task-switch prediction (Shen et al. 2006) and user business (Kapoor and Horvitz 2007), although some amount of explicit user feedback has been required at times (such as specifying the current user tasks). Our work focuses on building a good predictive model of the entire event log and extracting salient motifs in the data based on purely unsupervised methods.

6 Future Work

We presented a simple efficient learning technique to effectively aggregate prediction of multiple features when there are many possible classes to predict and in a nonstationary setting. For action prediction, we are investigating other tasks such as folder (directory) prediction on our own TaskTracer logs. For motif discovery, we are exploring algorithmic variations, as well as assessing utility in prediction and in user modeling.

Acknowledgments

Thanks to the reviewers, whose comments improved the presentation. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. FA8750-07-D-0185/0004. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the Air Force Research Laboratory.

References

- Bao, X.; Herlocker, J.; and Dietterich, T. 2006. Fewer clicks and less frustration: Reducing the cost of reaching the right folder. In *IUI*.
- Bekkerman, R., and McCallum, A. 2004. Automatic categorization of email into folders: benchmark experiments on enron and sri corpora. Technical report, Enron and SRI corpora (Technical Report IR-418). CIIR, UMass.
- Chambers, N., and Jurafsky, D. 2008. Unsupervised learning of narrative event chains. In *ACL*.
- Davison, B. D., and Hirsh, H. 1998. Predicting sequences of user actions. In *AAAI-98/ICML'98 Workshop on Predicting the Future: AI Approaches to Time Series Analysis*.

- Dragunov, A.; Dietterich, T.; Johnsrude, K.; McLaughlin, M.; Li, L.; and Herlocker, J. 2005. Tasktracer: a desktop environment to support multi-tasking knowledge workers. In *IUI*, 75–82. New York, NY, USA: ACM.
- Greenberg, S. 1988. Using Unix: collected traces of 168 users. Technical report, University of Calgary, Alberta.
- Horvitz, E. 1999. Principles of mixed-initiative user interfaces. In *ACM CHI*.
- Kapoor, A., and Horvitz, E. 2007. Principles of lifelong learning for predictive user modeling. In *Proc. Eleventh Conference on User Modeling (UM 2007)*.
- Korvemaker, B., and Greiner, R. 2000. Predicting UNIX command lines: Adjusting to user patterns. In *AAAI/IAAI*.
- Madani, O., and Connor, M. 2008. Large-scale many-class learning. In *SIAM Conference on Data Mining*.
- Madani, O., and Huang, J. 2008. On updates that constrain the number of connections of features during learning. In *ACM KDD*.
- Madani, O.; Bui, H.; and Yeh, E. 2009. Prediction and discovery of users' desktop behavior. Technical report, AI Center, SRI International.
- Manning, C. D., and Schütze, H. 1999. *Foundations of Statistical Natural Language Processing*. The MIT Press.
- Segal, R. B., and Kephart, J. O. 2000. Incremental learning in SwiftFile. In *ICML*.
- Shen, J.; Li, L.; Dietterich, T.; and Herlocker, J. 2006. A hybrid learning system for recognizing user tasks from desktop activities and email messages. In *IUI*.