Fast, Accurate, and Practical Identity Inference Using TV Remote Controls

Mariano Phielipp
Digital Home Group
Intel Corporation
Chandler, AZ
mariano.j.phielipp@intel.com

Magdiel Galan
Department of Computer Science and
Engineering
Arizona State University
magdiel.galan@asu.edu

Richard Lee
Technology and Manufacturing Group
Intel Corporation
Chandler, AZ
richard.m.lee@intel.com

Branislav Kveton
Intel Labs
Santa Clara, CA
branislav.kveton@intel.com

Jeffrey Hightower
Intel Labs
Seattle, WA
jeffrey.r.hightower@intel.com

Abstract

Non-invasive identity inference in the home environment is a very challenging problem. A practical solution to the problem could have far reaching implications in many industries, such as home entertainment. In this work, we consider the problem of identity inference using a TV remote control. In particular, we address two challenges that have so far prevented the work of Chang *et al.* (2009) from being applied in a home entertainment system. First, we show how to learn the patterns of TV remote controls incrementally and online. Second, we generalize our results to partially labeled data. To achieve our goal, we use state-of-the-art methods for max-margin learning and online convex programming. Our solution is efficient, runs in real time, and comes with theoretical guarantees. It performs well in practice and we demonstrate this on 4 datasets of 2 to 4 people.

Introduction

Providing multimedia content in a personalized TV environment that aligns the most with the interests of its consumers is a challenging problem for both service providers and content developers. This problem becomes even more challenging in families, where the recognition of individual members is highly desirable. The goal is to provide the best personalized experience for various multimedia contents, such as TV, on-demand programming, interactive media, targeted advertising, online gaming, and many others.

In our paper, we discuss practical challenges in building a non-invasive system for identifying TV viewers. The system minimizes the need of the TV viewers to log into their entertainment profiles, or being identified by an invasive method,

such as a camera combined with face recognition. Our solution is designed as follows. Whenever the TV viewer reveals identity, we use the corresponding remote control data, such as a sequence of button presses and accelerometer readings, to train the predictor of the person. When the identity of the viewer is unknown, we infer the identity based on the remote control data. Although our solution is specific to the remote control domain, note that many ideas in the paper generalize beyond it. In particular, our work is a prime example of how to turn an offline and fully-supervised solution into an online solution on partially labeled data.

The basis of our engineering efforts relies on the paper of Chang *et al.* (2009), which employs support vector machines (SVMs) and max-margin Markov networks (M^3Ns) (Taskar, Guestrin, and Koller 2004) to infer the identify of TV remote control users. We address two main problems that had so far prevented this work from being implemented in a home entertainment system. First, the original methodology assumes completely labeled data. Unfortunately, TV viewers usually provide very little feedback about their identity. Second, the approach of Chang *et al.* (2009) does not allow for an incremental improvement of learned predictors. This is necessary since remote control data are usually unavailable in advance, and only become sporadically available as time progresses.

Considering our minimal invasive system setup, in which we rarely observe labeled data, we show that there is simply not enough data to build a reasonably good manifold, which can be used by semi-supervised learning algorithms. In turn, we focus on supervised learning only and try to evaluate how many labeled examples are needed to learn good predictors. First, we answer this question in the offline setting. Second, we show how to learn online from completely labeled data.

Household	Participants	Sessions
1	4	458
2	2	124
3	3	28
4	2	90
5	4	340

Table 1: Households statistics.

Finally, we relax the assumption on completely labeled data. One of our results is that the identity of remote control users can be inferred with an acceptably high accuracy even when only 20 percent of data are labeled.

These results are obtained using state-of-the-art methods for online convex programming (Zinkevich 2003). Although our solution is simple and learned online, the accuracy of the solution is often comparable to Chang *et al.* (2009). To further improve the solution, we propose a new way of training max-margin Markov networks online (Ratliff, Bagnell, and Zinkevich 2007). The algorithm runs in real time and can be implemented on a commercial platform. Finally, both of our solutions are comprehensively evaluated on 4 remote control datasets of 2 to 4 people, and compared to online and offline majority class baselines.

The following notation is used in the paper. The symbols $\mathbf{x}_t$ and $y_t \in \{-1, 1\}$ denote the t-th data point and its label, respectively. The data points $\mathbf{x}_t$ are divided into labeled and unlabeled sets, l and u, and the labels y_t are observed for the labeled set only. The cardinality of the labeled and unlabeled sets is $n_l = |l|$ and $n_u = |u|$, respectively, and $T = n_l + n_u$ denotes the total number of training examples.

Remote control dataset

Our data set consists of data collected on five households for a period of one to three weeks in which the number of users for the households varied between two and four. This is the same dataset featured in the Chang *et al.* (2009) paper. Each household had a data collection system that consisted of a tri-axis accelerometer attached to a TV remote control, the corresponding accelerometer receiver, a universal infrared receiver to capture button presses from the remote control, and a laptop to which the receivers were connected. The laptop logged and time stamped the data from the sensors using 100 nanoseconds resolution.

For our analysis, all data of interest revolved around the remote control activity in the form of button press selections. Since we seek to associate combined accelerometer and button readings to individual users, we concern ourselves with the behaviors just before and just after each of the button presses. To study just how much before and how much after, we implement four different capture windows at 0.5, 1, 2, 4 seconds with respect to the button press. The windows were designed to help capture the hand motions preceding, centered, and succeeding each button press. We believe the variety in window sizes is sufficient to capture the uniqueness in hand motions for each of the users.

Since data was collected using 100 nanoseconds resolution, the data contained within each of those windows

were used to generate the features that defined each button press instance. Some of the features associated with the accelerometer data include: energy, fundamental frequency, range, mean, and variance for each of the axes, as well as correlation among each pair of axes. Some of the features associated with the infrared button signal from the remote control include the button code, press duration and number of times the button code was sequentially transmitted.

In all, a total of 372 combined features defined each button press instance. The time stamped instances were then group into sessions. A session represents the periods of time in which there is continuous remote control activity. A session ends when it is determined that the remote control still idle. Table 1 illustrates the total number of sessions for each household. Notice that because of the scarcity of sessions, Household 3 was not considered in this study.

For our experiments, all the instances corresponding to the same session are aggregated into a session-level instance representation as the mean for the session. Button that were rarely use were discarded before aggregation. We then normalize with respect to all sessions.

Algorithms

This section reviews online and offline learning algorithms, which are used in the experimental section.

Supervised max-margin learning

Support vector machines (SVMs) (Vapnik 1995) are a standard algorithm for learning max-margin discriminators. The learning problem is formulated as:

$$\min_{f \in \mathcal{H}_K} \frac{1}{T} \sum_{t=1}^{T} V(f, \mathbf{x}_t, y_t) + \gamma \|f\|_K^2, \qquad (1)$$

where $V(f, \mathbf{x}, y) = \max\{1 - yf(\mathbf{x}), 0\}$ represents the *hinge loss*, f is a function from a *reproducing kernel Hilbert space (RKHS)* $\mathcal{H}_K$, and $\|\cdot\|_K$ is the RKHS norm that measures the complexity of f. The tradeoff between the regularization of f and minimizing the losses $V(f, \mathbf{x}_t, y_t)$ is controlled by the parameter γ. In all experiments, $\gamma = 0.01$.

The temporal structure of our problem allows for improving SVM predictions (Chang, Hightower, and Kveton 2009). Chang *et al.* (2009) applied max-margin Markov networks to capture this structure and reported significant improvements over SVMs. In this work, we propose a much simpler solution. The predictions $\hat{y}_t = \text{sgn}(f(\mathbf{x}_t))$ are smoothed out over time as:

$$(\tilde{y}_1, \ldots, \tilde{y}_T) = \arg \max_{k_1, \ldots, k_T}$$

$$\left\{ \phi \sum_{i=1}^{T-1} \mathbb{1}\{k_i = k_{i+1}\} + \right.$$

$$\left. \sum_{i=1}^{T} \mathbb{1}\{\text{sgn}(f(\mathbf{x}_i)) = k_i\} |f(\mathbf{x}_i)| \right\}, \quad (2)$$

where $\mathbb{1}\{\cdot\}$ is an indicator and ϕ determines the importance of label smoothing. The higher the value of the parameter ϕ,
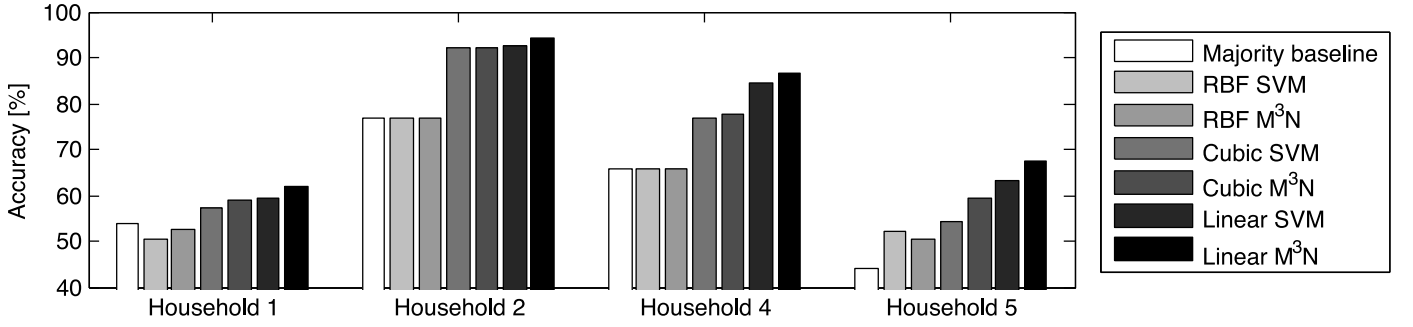
Figure 1: Comparison of 7 solutions on 4 remote control datasets. From left to right, we report the accuracy of the majority class baseline, RBF SVMs, RBF M^{3}Ns, cubic SVMs, cubic M^{3}Ns, linear SVMs, and linear M^{3}Ns.

the more the labels $\tilde{y}_t$ are smoothed over time. When $\phi = 0$, the predictions $\tilde{y}_1, \ldots, \tilde{y}_T$ are identical to the predictions of the corresponding SVM. In all experiments, $\phi = 1$.

When compared to M^{3}Ns, the smoothed-out predictor has one significant advantage. It is much easier to learn because it reuses the SVM decision boundary and the only additional parameter is ϕ. Due to its simplicity, the predictor is unlikely to be as good as the M^{3}Ns of Chang *et al.* (2009). However, note that it models the same kind of dependencies. Hence, in the rest of the paper, we refer to it informally as an M^{3}N.

Semi-supervised max-margin learning

Manifold regularization of SVMs (Belkin, Niyogi, and Sindhwani 2006) is one way of combining max-margin learning and semi-supervised learning on graphs. This learning problem is formulated as:

$$\min_{f \in \mathcal{H}_K} \frac{1}{n_l} \sum_{t \in l} V(f, \mathbf{x}_t, y_t) + \gamma \|f\|_K^2 + \gamma_u \mathbf{f}^\top L \mathbf{f}, \quad (3)$$

where $\mathbf{f} = (f(\mathbf{x}_1), \ldots, f(\mathbf{x}_T))$ and L denotes the Laplacian of the data adjacency graph, which is represented by a matrix W of pairwise similarities w_{ij}. The similarities w_{ij} are often computed as:

$$w_{ij} = \exp[- \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / (2K\sigma^2)], \quad (4)$$

where K is the number of features and the heat parameter σ denotes the mean of their standard deviations. This approach is adopted in the rest of the paper.

The scalar γ_u controls the importance of unlabeled examples. When $\gamma_u = 0$, the objective of manifold regularization (3) is identical to the objective of SVMs (1).

Online max-margin learning

Online learning of SVMs (1) can be formulated as an online convex programming problem.

Online convex programming (Zinkevich 2003) involves a convex feasible set $\mathcal{F} \subset \mathbb{R}^K$ and a sequence of convex functions $c_t : \mathcal{F} \to \mathbb{R}$. At each time step t, we choose an action $f_t \in \mathcal{F}$ based on the past functions $c_1, \ldots, c_{t-1}$ and actions $f_1, \ldots, f_{t-1}$, and the goal is to minimize the regret:

$$\sum_{t=1}^{T} c_t(f_t) - \min_{f \in \mathcal{F}} \sum_{t=1}^{T} c_t(f). \quad (5)$$

The above regret can be minimized on the order of $O(\sqrt{T})$ by the gradient update:

$$f_{t+1} = P(f_t + \eta \nabla c_t(f_t)), \quad (6)$$

where $\eta = \sqrt{T}$ is a learning rate, $\nabla c_t(f_t)$ is the gradient of the function c_t at the point f_t, and $P(\cdot)$ is a projection to the feasible set $\mathcal{F}$. When the cost function is defined as:

$$c_t(f) = V(f, \mathbf{x}_t, y_t) + \gamma \|f\|_K^2, \quad (7)$$

our gradient update minimizes the regret with respect to the objective of SVMs. For linear SVMs, the cost function simplifies to:

$$c_t(f) = V(f, \mathbf{x}_t, y_t) + \gamma \|f\|_2^2. \quad (8)$$

Similarly to Equation 2, the online SVM predictor can be smoothed out over time as:

$$(\tilde{y}_1, \ldots, \tilde{y}_t) = \arg \max_{k_1, \ldots, k_t}$$
$$\left\{ \phi \sum_{i=1}^{t-1} \mathbb{1}\{k_i = k_{i+1}\} + \right.$$
$$\left. \sum_{i=1}^{t} \mathbb{1}\{\mathrm{sgn}(f(\mathbf{x}_i)) = k_i\} |f(\mathbf{x}_i)| \right\}, \quad (9)$$

where $\mathbb{1}\{\cdot\}$ is an indicator and ϕ determines the importance of label smoothing. In practice, the maximization (9) should be performed over a window of most recent examples rather than the entire history. In our experiments, we limit the window to 10 most recent examples. We observed no significant difference in results when the size of the window varies from 5 to 20 examples.

Experiments

Our experiments are divided into 4 groups. In each of them, we gradually relax the assumptions of Chang *et al.* (2009) on learning offline and from completely labeled data. Naturally, we progressively consider more and more practical solutions to identity inference using TV remote controls.

All experiments are done in MATLAB. Manifold regularization of SVMs is evaluated based on the implementation of Belkin *et al.* (2006). Offline learning of SVMs is performed using LIBSVM (Chang and Lin 2001). All offline learning results are obtained by 10-fold cross-validation.

Dataset	L	Majority baseline	SVM	Manifold regularization of SVMs								
				10^{-4}	10^{-3}	10^{-2}	10^{-1}	1	10^1	10^2	10^3	10^4
	20	76.61	88.71	71.77	73.39	78.23	80.65	76.61	83.87	85.48	85.48	83.87
	40	76.61	92.74	56.45	56.45	70.16	83.87	82.26	83.87	77.42	76.61	77.42
Household 2	60	76.61	93.55	41.94	41.94	37.90	84.68	82.26	74.19	66.13	70.97	68.55
	80	76.61	91.13	82.26	83.06	86.29	89.52	79.84	77.42	75.00	74.19	77.42
	100	76.61	92.74	66.13	67.74	75.00	91.13	83.06	76.61	69.35	33.87	40.32
	20	65.56	68.89	34.44	34.44	34.44	34.44	34.44	34.44	34.44	34.44	33.33
	40	65.56	68.89	34.44	34.44	34.44	34.44	34.44	34.44	34.44	34.44	34.44
Household 4	60	65.56	70.00	34.44	34.44	34.44	35.56	34.44	36.67	37.78	34.44	34.44
	80	65.56	81.11	40.00	40.00	38.89	37.78	34.44	34.44	34.44	34.44	34.44
	100	65.56	84.44	50.00	50.00	51.11	50.00	34.44	34.44	34.44	34.44	34.44

Table 2: Comparison of SVMs, manifold regularization of SVMs, and the majority class baseline on 2 remote control datasets. Manifold regularization of SVMs is performed for various regularization parameters $\gamma_u \in [10^{-4}, 10^4]\gamma$. The fraction of labeled examples L varies from 20 to 100 percent.
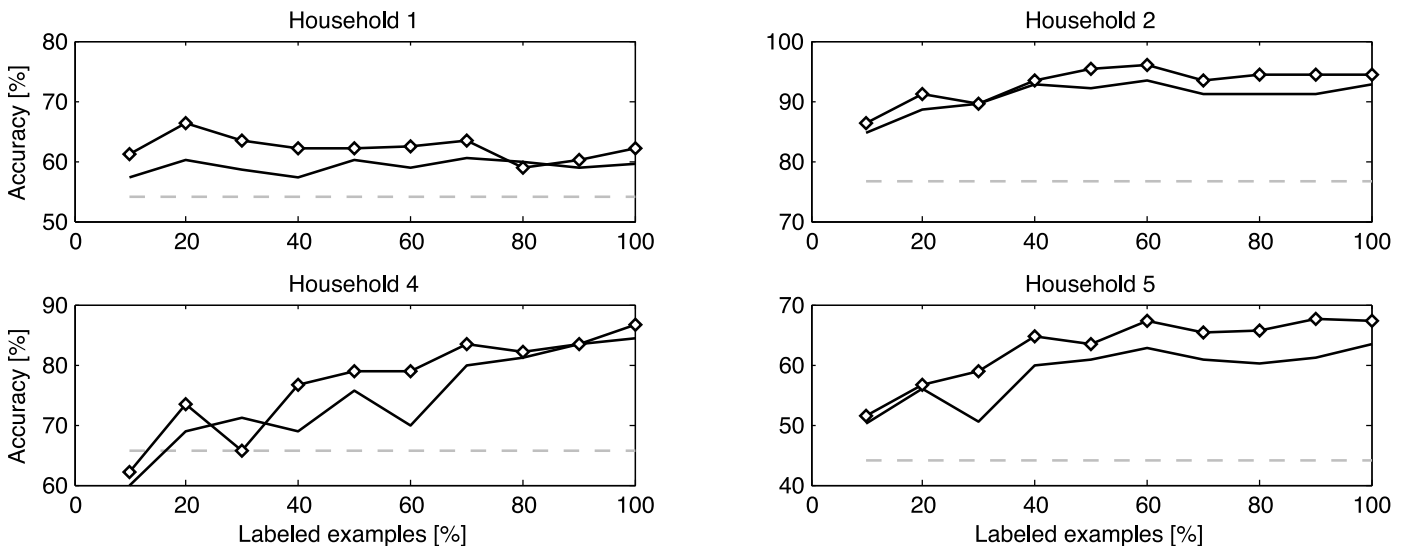

Figure 2: Comparison of SVMs (black lines) and M³Ns (black lines with diamonds) on 4 remote control datasets. The methods are compared by their accuracy, which is reported as a function of the fraction of labeled examples. We also report the accuracy of the majority class baseline (dashed gray lines).

Complex decision boundaries

The first experiment evaluates the benefit of using non-linear predictors in our domain. Moreover, we try to improve these predictors by utilizing the structure of our problem. In short, we compare 6 max-margin classifiers: SVMs (1) with linear, cubic, and RBF kernels; and M³Ns (2), which are computed over the same set of kernels.

Our results are shown in Figure 1. Based on these results, we conclude that the simplest decision boundaries yield the highest accuracy on all datasets. Therefore, learning of non-linear decision boundaries is not beneficial at all. A possible cause for this result is that the dimensionality of our datasets is too high in comparison to the number of training examples (Table 1). In turn, linear decision boundaries yield close-to-optimal results. As a result, the rest of our experiments focus on linear models only.

In addition, note that M³Ns yield better results than SVMs

for almost all households. The benefit of using the temporal structure is lower than reported by Chang *et al.* (2009) since we use a simpler model.

Partially labeled data

The second experiment is focused on learning from partially labeled data. Our main objective is to evaluate the feasibility of learning in this setting, and explore the utility of unlabeled data.

As a representative method for semi-supervised learning, we select manifold regularization of SVMs (3). The method is evaluated on 2 of our datasets and our results are reported in Table 2.[1] Based on these results, we conclude that semi-supervised learning is not suitable for our domain. Manifold

---

[1]Manifold regularization of SVMs cannot be straightforwardly generalized to multi-class problems. Therefore, the method is evaluated only on households of 2 people.
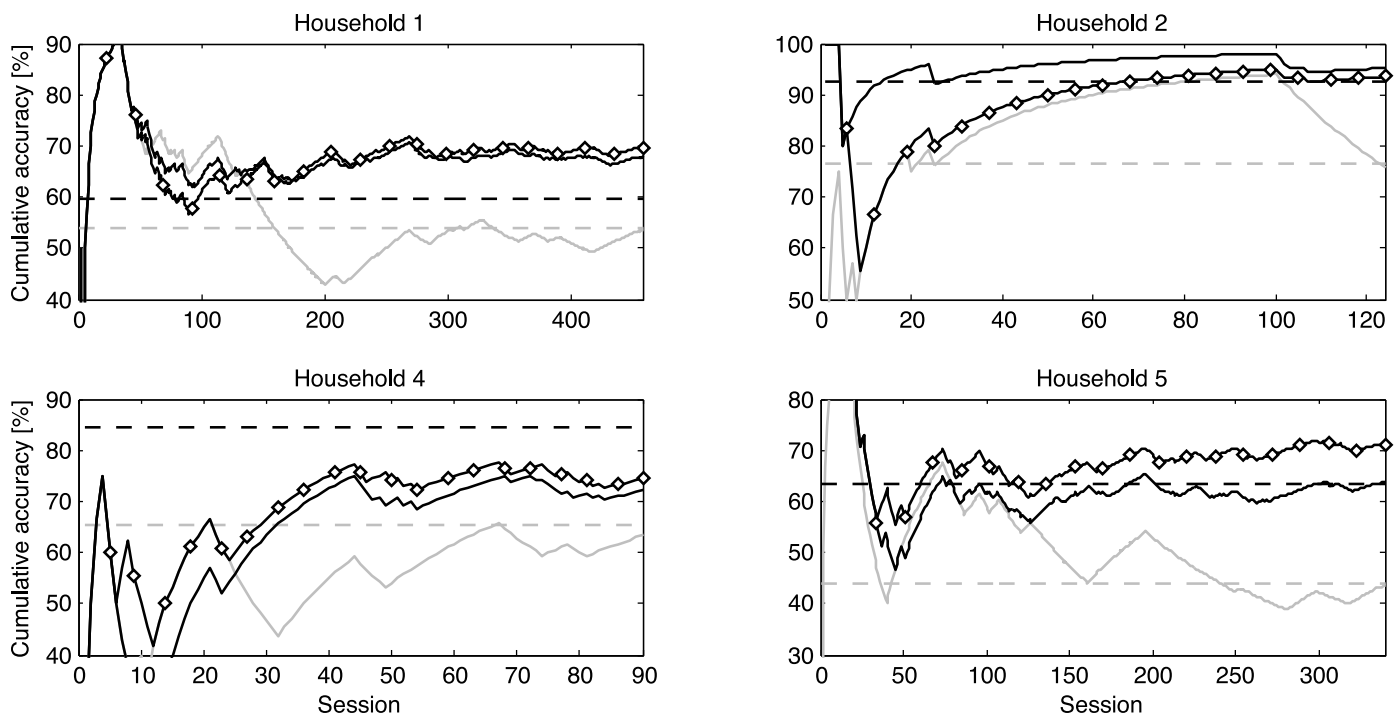
Figure 3: Comparison of online-learned SVMs (black lines) and M³Ns (black lines with diamonds) on 4 remote control datasets. The methods are compared by their cumulative accuracy up to each time step. We also report the accuracy of online and offline majority class baselines (solid and dashed gray lines), and offline-learned SVMs (dashed black lines).

regularization of SVMs is often worse than supervised learning of SVMs, and in many cases, it does not even outperform the majority class baseline. This trend is likely caused by the lack of a manifold in our data.

Since the unlabeled data does not seem very valuable, we disregard the data in the rest of our experiments and focus on supervised learning. Figure 2 shows that supervised learning with linear SVMs yields extremely good results. Even when a small portion of data is labeled, such as 20 percent, SVMs beat the majority class baseline. Smoothing of the results by M³Ns yields an additional improvement in accuracy.

Online learning

Up to this point, we studied a version of our problem, where data are collected in advance. However, in practice, the data are generated on-the-fly whenever someone uses the remote control. A practical solution to our problem should take this property into account. It should also improve over time and adapt when the patterns of users change. One way of obtaining this solution is by formulating our problem as an online learning problem (Cesa-Bianchi and Lugosi 2006).

In this work, our online learning algorithm iteratively descends the gradient (Equation 6) of the cost function (Equation 8), and minimizes the regret with respect to the optimal solution to linear SVMs. To evaluate our algorithm, we emulate the environment where training examples are introduced over time. All examples are labeled and shown one at a time. Our learner is compared to offline-learned SVMs, the majority class baseline, and an online majority class baseline. The

online baseline votes for the most frequent label in retrospect (Blum 1996).

Our results are shown in Figure 3. Based on these results, we conclude that online learning of linear SVMs is a viable solution to our problem. The algorithm quickly outperforms the majority class baseline and once learning is completed, it performs as well as offline-learned SVMs in 3 out of 4 cases. Moreover, our solution also outperforms the online majority class baseline. Thus, we may conclude that the combined set of button-press and accelerometer features provides as good or better classifier than the majority vote in retrospect. As in our previous experiments, smoothing through M³Ns usually improves the SVM results.

Online learning with partially labeled data

The last experiment essentially combines the ideas of online learning with learning from partially labeled data. This setup is relevant in our problem because remote control users only rarely provide feedback about their identity, such as logging into their profile on a TV or a home entertainment system.

In this experiment, we assume that only a fraction of data is labeled. Otherwise, our experimental setup is identical to the last experiment. The fraction of labeled data varies from 10 to 100 percent in 10 percent increments, and the data are chosen uniformly from the entire dataset.

Our results are shown in Figure 4. Based on these results, we conclude that online learning of linear SVMs is a viable solution to predicting the identity of remote control users. In particular, even when a small portion of data is labeled, such
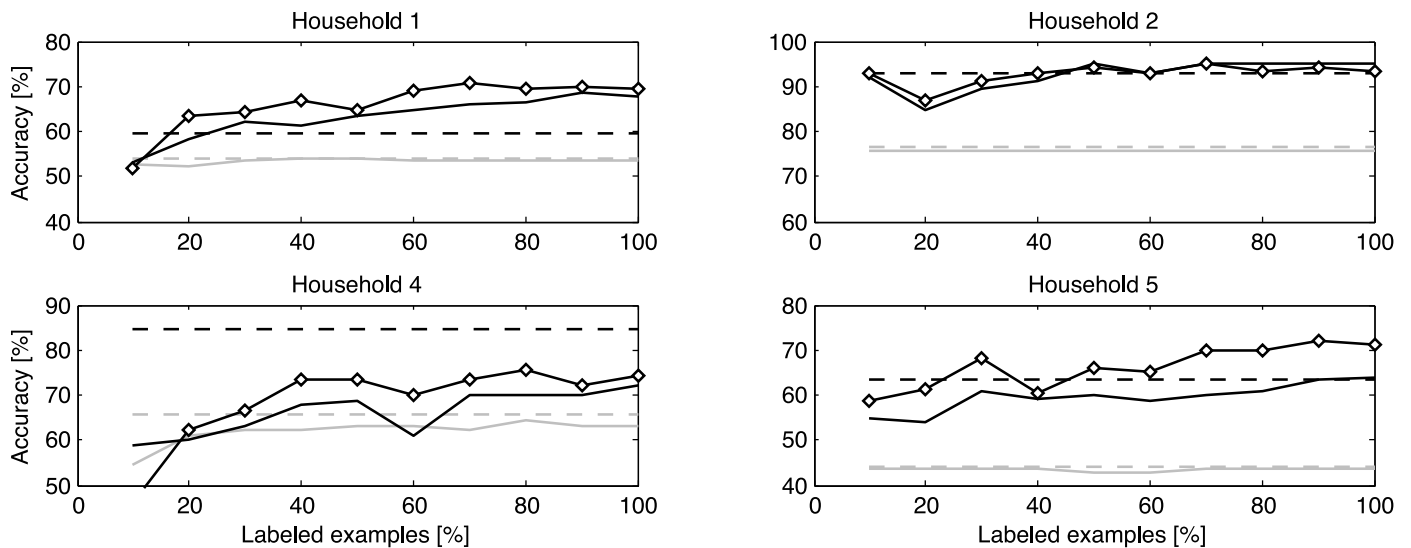
Figure 4: Comparison of online-learned SVMs (black lines) and M³Ns (black lines with diamonds) on 4 remote control datasets. The methods are compared by their accuracy, which is reported as a function of the fraction of labeled examples. We also report the accuracy of online and offline majority class baselines (solid and dashed gray lines), and offline-learned SVMs (dashed black lines). The offline-learned SVMs are trained on completely labeled data.

as 20 percent, online-learned SVMs outperform both offline and online majority class baselines in 3 out of 4 cases. As in our previous experiments, smoothing through M³Ns usually improves the SVM results. Finally, note that online-learned M³Ns sometimes perform better than offline-learned SVMs.

Conclusions

Non-invasive identity inference in the home environment is a challenging problem. In our paper, we build on the work of Chang *et al.* (2009) and show how to make it more practical. In particular, we propose an online algorithm that learns the identity of remote control users over time and from partially labeled data. The algorithm runs in real time, performs well in practice, and comes with theoretical guarantees on its performance.

The accuracy of our predictor is nowhere near to 100 percent. This is not surprising due to the non-invasive character of our sensor. On the other hand, this precludes our approach from being applied in domains, where users have to be identified with a high precision or they are immediately annoyed. One of the suitable domains for our approach is targeted TV advertising. In this context, ads are already present, and their personalization would likely only improve the quality of the ads.

Our approach relies on partially labeled data, which is one of its shortcomings. This issue is less significant than it may seem due to recent changes in the consumer electronics business. In particular, many new TVs can be used to access the Internet. To gain the access, consumers are often required to create their profiles. The profiles can be viewed as a labeling mechanism, which labels our data whenever a user logs into the profile.

In our future work, we plan to use non-remote control pat-

terns, like surfing and viewing habits, to improve our predictor. Also, we would like to take in account semantically relevant time periods, such as weekends, holidays, and seasonal changes in programming. Finally, we plan to collect a larger dataset and investigate the value of semi-supervised learning in this dataset.

References

Belkin, M.; Niyogi, P.; and Sindhwani, V. 2006. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research* 7:2399–2434.

Blum, A. 1996. On-line algorithms in machine learning. In *Online Algorithms*, volume 1442 of *Lecture Notes in Computer Science*, 306–325. Springer.

Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, Learning, and Games*. New York, NY: Cambridge University Press.

Chang, C.-C., and Lin, C.-J. 2001. *LIBSVM: a library for support vector machines*. Software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.

Chang, K.; Hightower, J.; and Kveton, B. 2009. Inferring identity using accelerometers in television remote controls. In *Proceedings of the 7th International Conference on Pervasive Computing*, 151–167.

Ratliff, N.; Bagnell, A.; and Zinkevich, M. 2007. (Online) subgradient methods for structured prediction. In *Proceedings of the 11th International Conference on Artificial Intelligence and Statistics*.

Taskar, B.; Guestrin, C.; and Koller, D. 2004. Max-margin Markov networks. In *Advances in Neural Information Processing Systems 16*.

Vapnik, V. 1995. *The Nature of Statistical Learning Theory*. New York, NY: Springer-Verlag.

Zinkevich, M. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, 928–936.