

Learning User Profile from Traces

Ugo Galassi, Attilio Giordana and Dino Mendola
Dipartimento di Informatica, Università del Piemonte Orientale
Spalto Marengo 33, Alessandria, Italy

Abstract

This paper presents a method for automatically constructing a sophisticated user profile from traces of user behavior. User profile is encoded by means of a Hierarchical Hidden Markov Model (HHMM). The HHMM is a well formalized tool suitable to model complex patterns in long temporal or spatial sequences. The method described here is based on a recent algorithm, which is able to synthesize the HHMM structure from a set of logs of the user activity. The algorithm follows a bottom-up strategy, in which elementary facts in the sequences (motives) are progressively grouped, thus building the abstraction hierarchy of a HHMM, layer after layer. The induction strategy has been designed in order to deal with events characterized by a sparse structure, where gaps filled by irrelevant facts can be intermixed with the relevant ones. The method is firstly evaluated on artificial data. Then a user identification task, from real data, has been considered. A first experiment with a set of 14 different users produced encouraging results.

1. Introduction

Building profiles for processes and for interactive users, is a important task in intrusion detection. This paper presents an algorithm, based on Hierarchical Hidden Markov Model [5], for learning sophisticated profiles (models) of the behavior of processes. The algorithm discovers typical "motives" of a process behavior, and correlates them into a hierarchical model. Motives can be interleaved with possibly long gaps where no regular behavior is detectable. We assume that motives could be affected by noise due to non-deterministic causes. Noise is modeled as insertion, deletion and substitution errors according to a common practice followed in Pattern Recognition. An approach to deal with such kind of patterns, which reported impressive records of successes in speech recognition [9] and DNA analysis [4], is the one based on Hidden Markov Model (HMM) [8]). However, applying HMM does not reduce to simply running a learning

algorithm but it require to spend a considerable effort to individuate a suitable structure for the HMM. Moreover, complex applications require to construct an ad hoc system, where several partial HMMs are developed and integrated with procedural knowledge obtained from experts of the domain. A formal framework to design and train complex HMMs is represented by the Hierarchical Hidden Markov Model (HHMM) [5]. The problem of estimation HMM and HHMM parameters has been widely investigated while little has been done in order to learn their structure. A few proposals can be found in the literature in order to learn the structure of HMM. A novelty, in this sense, is represented by a recent paper by Botta et Al. [2], which proposes a method for automatically inferring from sequences, and possibly domain knowledge, both the structure and the parameters of complex HHMMs.

In this paper, the learning algorithm proposed in [2] is briefly overviewed and is experimentally evaluated on two profiling case studies. The first one consists of a suite of artificial traces automatically generated by a set of given HHMMs. The challenge for the algorithm is to reconstruct the original model from the traces. It will be shown that the algorithm is able to learn HHMMs very similar to the original ones, in presence of noise and distractors.

The second case study refers to the problem of constructing a discriminative model for a user typing on a keyboard [1, 3, 6]. The results reported with a set of 20 different users are encouraging.

2. The Hierarchical Hidden Markov Model

A Hierarchical Hidden Markov Model is a generalization of the Hidden Markov Model, which is a stochastic finite state automaton [8] defined by a tuple $\langle S, O, A, B, \pi \rangle$, where:

- S is a set of states, and O is a set of atomic events (observations),
- A is a probability distribution governing the transitions from one state to another. Specifically, any member

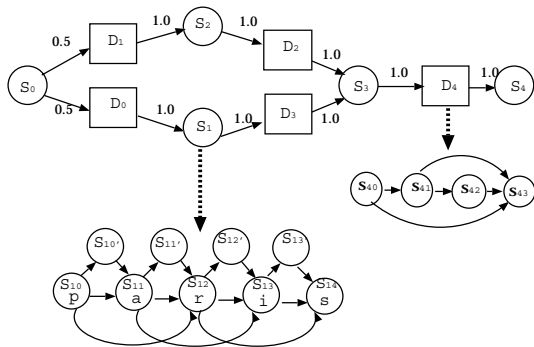


Figure 1. Example of Hierarchical Hidden Markov Model. Circles denotes states with observable emission, whereas rectangles denote gaps.

$a_{i,j}$ of A defines the probability of the transition from state s_i to state s_j , given s_i .

- B is a probability distribution governing the emission of observable events depending on the state. Specifically, an item $b_{i,j}$ belonging to B defines the probability of producing event O_j when the automaton is in state s_i .
- π is a distribution on S defining, for every $q_i \in S$, the probability that s_i is the initial state of the automaton.

A difficulty, related to a HMM defined in this way, is that, when the set of states S grows large, the number of parameters to estimate (A and B) rapidly becomes intractable.

A second difficulty is that the probability of a sequence being generated by a given HMM decreases exponentially with its length. Then, complex and sparse events become difficult to discover.

The HHMM proposed by Fine, Singer and Tishby [5] is an answer to both problems. On one hand, the number of parameters to estimate is strongly reduced by assigning a null probability to many transitions in distribution A , and to many observations in distribution B . On the other hand, it allows a possibly long chain of elementary events to be abstracted into a single event, which can be handled as a single item. This is obtained by exploiting the regular languages property of being closed under substitution, which allows a large finite state automaton to be transformed into a hierarchy of simpler ones.

More specifically, numbering the hierarchy levels with ordinals increasing from the highest towards the lowest level, observations generated in a state s_i^k by a stochastic automaton at level k are sequences generated by an automaton at level $k + 1$. Moreover, no direct transition may occur between the states of different automata in the hierarchy. An example of HHMM is given in Figure 1.

The advantage of the hierarchical structure, as defined by [5], may help very much the inference of the entire structure of the automaton by part of an induction algorithm.

The research efforts about HHMM mostly concentrate on the algorithms for estimating the probabilities governing the emissions and the transition from state to state. In the seminal paper by Fine et al. [5], the classical Baum-Welch algorithm is extended to the HHMM. In a more recent work, Murphy and Paskin [7] derive a linear (approximated) algorithm by mapping a HHMM into a Dynamic Bayesian Network.

3. Learning Algorithm Overview

The basic algorithm [2] is bottom-up and constructs the HHMM hierarchy starting from the lowest level. The first step consists in searching for possible motives, i.e., short chains of consecutive symbols that appear frequently in the learning traces, and building a HMM for each one of them. As motives are considered independently one from another, this phase tends to produce also models for spurious motives, which in a second time should be discarded. At the same time, it may happen that relevant motives be disregarded just because their frequency is not high enough. However, such kind of errors will be fixed at a second time. The HMMs learned so far, are then used as feature constructors. Each HMM is labeled with a different *name* and the original sequences are rewritten into the new alphabet defined by the set of names given to the models. Every subsequence, which can be attributed to a specific HMM, is replaced by the corresponding name. We will discuss in Section 3 how possible conflicts are solved.

The subsequences between two motives, not attributed to any model, are considered gaps and will be handled by means of special construct called *gap*. We will call this last operation *sequence abstraction*. After this basic cycle has been completed, an analogous procedure is repeated on the abstracted sequences. Models are now built for sequences of *episodes*, searching for long range regularities among co-occurrent motives. In this process, spurious motives not showing significant regularities can be discarded. The major difference, with respect to the first learning step, is that the models built from the abstract sequences, are now observable markov models. This makes the task easier and decreases the computational complexity. In this step, models (*gaps*) are built also for the long intervals falling between consecutive motives.

In principle, the abstraction step could be repeated again to the learning sequences, building a third level of the hierarchy, and so on. However, up to now, we considered only problems where two levels are sufficient.

After building the HHMM structure in this way, it can be refined using standard training algorithms like the ones

proposed in [5, 7]. However, two other refinement methods are possible.

The first method concerns the recovery of episodes that have been lost in the primary learning phase because they did not have a sufficient statistical evidence. As said above, this missed information has actually been modeled by *gaps*. A nice property of the HHMM is that sub-models in the hierarchy have a loose interaction with one another, and so their structure can be reshaped without destroying the global structure. This means that the model of a gap can be transformed into the model of a motif later on, when further data will be available. This is actually done on demand: all sub-sequences attributed to a given *gap* are collected, creating a new learning set where the learning process is repeated. If there is now evidence for a motif, a model is built up and replaced to the *gap*.

The second method consists in repeating the entire learning cycle using as learning set only the portion of the sequences where the instance of the previously learned HHMM have been found with sufficient evidence. Repeating the procedure allows more precise models to be learned for motives, because false motives will no longer participate to the learning procedure.

4. Evaluation on Artificial Traces

A specific testing procedure has been designed in order to monitor the capability of the algorithm of discovering "known patterns" hidden in trace artificially generated by a handcrafted HHMMs. Random noise and spurious motives have been added to all sequences filling the gaps between consecutive motives.

Three target HHMMs, each one constructed according to a two level hierarchy have been used to generate a set of 72 learning tasks (24 for every model). Every learning task consists of a set of 330 traces. The 90% of the sequences contain an instance of a target HHMM that should be discovered by the learning program, whereas the 10% contain sequences of spurious motives non generated by the target HHMM. The sequence length ranges from 80 to 120 characters.

The structure for the high level of the three models is shown in Figure 2. Every state at the high level emits a string (motif) generated by an HMM at the low level, indicated with a capital letter (A,B,C,D,E). Another HMM (F), has been used to generate spurious motives. The gaps between motives have been filled with subsequences containing random noise.

The evaluation of the obtained results has been done using the notion of *bayes classification error* between two (or more HHMMs). Formally, given two HHMMs, λ_1 and λ_2 , and the set L of all possible traces, which can be generated




Figure 2. HHMM used for evaluation on artificial data

by λ_1 or λ_2 , the Bayes classification error $C(\lambda_1, \lambda_2)$ is defined as:

$$C(\lambda_1, \lambda_2) = \sum_{x \in L} [\min(p(\lambda_1|x), p(\lambda_2|x))]p(x) \quad (1)$$

being $p(\lambda_1|x)$ and $p(\lambda_2|x)$ the probability that, given a trace x , it has been generated by λ_1 or λ_2 , respectively, and $p(x)$ the a priori probability of x . We notice that the upper-bound for $C(\lambda_1, \lambda_2)$ is 0.5, when λ_1 and λ_2 are identical. In general, for N models, the upper-bound is given by the expression $1 - 1/N$.

In general, expression (1) cannot be computed because L is too large. Therefore, we adopted an approximate evaluation made using a subset of L stochastically sampled.

The bayes classification error 1 intervenes in the evaluation procedure in two different ways. A first way is to measure the quality of the learned models. A perfect learner should learn a model identical to the one used to generate the traces. Therefore, a learned model has to be considered as much accurate as much close to 0.5 the classification error, between it and the original model, is.

The second way is to estimate the difficulty of the learning task. It is reasonable to assume that the difficulty of identifying a model hidden in a set of traces grows along with the similarity among the motives belonging to the model and the spurious motives. Moreover, the difficulty grows also when the motives belonging to a same model become similar each other, because it becomes more difficult to discover the correspondence between a motif and the hidden state it has been emitted from. Therefore, the experimentation has been run using different versions of models A, B, C, D, E, F with different bayes classification error among them.

The results obtained under three different conditions of difficulty are summarized in Table 1. The similarity between the six kinds of motives has been varied from 0.2 to

Motives	0.2	0.4	0.55
Model (a)	0.48	0.46	0.45
Model (b)	0.47	0.42	0.42
Model (c)	0.43	0.42	0.41

Table 1. Bayes classification error between the target model and the learned model, versus the confusion among the basic motives (reported in the first line).

0.55. For every setting, the experiment has been repeated 8 times for each one of the three models. The reported results are the average over the 8 runs. In all cases, the bayes classification error has been estimated using a set of traces obtained by collecting 500 sequences generated from each one of the models involved in the specific comparison.

It appears that the performances suffer very little from the similarity among the motif models, and in all cases, the similarity between the original model and the learned model is very high ($C(\lambda_1, \lambda_2)$, is close to 0.5).

5. User Identification

The task consists in learning to identify a user from the its typing style on a keyboard. The basic assumption is that every user has a different way of typing that becomes particularly evident when he types words which are specifically important for him, such its name or words referring to his job. We considered such words as motives trying to build up a HHMM charactering a user on the basis of such motives.

A group of 14 users have been asked to type a sentence of 21 syllables on the same keyboard. A transparent program recorded, for every typed key, the duration and the delay two consecutive strikes creating a trace for every typing session. Each user provided ten repetitions of the sentence. Then, a dataset of 140 traces has been obtained, which has been partitioned into a learning set of 98 traces and a testing set of 42 traces. According to a standard procedure in machine learning, 14 HHMMs have been learned from the learning set. Then, the 14 HHMMs have been used to classify the traces in the testing set, according to the following procedure. For each HHMM M and for each trace t the forward-backward algorithm [8] is applied in order to estimate the probability for t of being generated by M . Then, t is attributed to the HHMM that has shows the highest probability. If such a HHMM is the model of the the user that has generated t , the classification is considered correct. Otherwise it is counted as a misclassification error. However, it may happen that all HHMMs show a null probability when t does not belong the language of anyone of them. This is

considered a rejection error. In, the specific case, 28 traces (66%) have been correctly classified with a very good margin (a strong difference between the probability assigned by the correct model and the other ones). The remaining 14 traces (33%) have been rejected.

6. Conclusion

We have proposed a method for automatically synthesize from traces (logs) profiles based on HHMMs. In preliminary tests on artificial datasets, the method succeeded in reconstructing two level HHMMs, whereas the results obtained on a task of user identification are encouraging. In fact, the high rejection rate simply means that the number of learning examples for each user was too small, so that some structural knowledge was missing from the HHMMs. Increasing the number of examples, the performances should increase. On the other hand, we consider very promising the fact that the distance between the models of the different users is very large.

7. Acknowledgments

The present work has been supported by the FIRB Project: WebMinds

References

- [1] S. Bleha, C. Slivinsky, and B. Hussein. Computer-access security systems using keystroke dynamics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-12(12):1217–1222, 1990.
- [2] M. Botta, U. Galassi, and A. Giordana. Learning complex and sparse events in long sequences. In *Proceedings of the European Conference on Artificial Intelligence, ECAI-04*, Valencia, Spain, August 2004.
- [3] M. Brown and S. Rogers. User identification via keystroke characteristics of typed names using neural networks. *International Journal of Man-Machine Studies*, 39:999–1014, 1993.
- [4] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison. *Biological sequence analysis*. Cambridge University Press, 1998.
- [5] S. Fine, Y. Singer, and N. Tishby. The hierarchical hidden markov model: Analysis and applications. *Machine Learning*, 32:41–62, 1998.
- [6] R. Joyce and G. Gupta. User authorization based on keystroke latencies. *Communications of the ACM*, 33(2):168–176, 1990.
- [7] K. Murphy and M. Paskin. Linear time inference in hierarchical hmms. In *Advances in Neural Information Processing Systems (NIPS-01)*, volume 14, 2001.
- [8] L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of IEEE*, 77(2):257–286, 1989.
- [9] L. Rabiner and B. Juang. *Fundamentals of Speech Recognition*. Prentice Hall, Englewood Cliffs, NY, 1993.