CERIAS

The Center for Education and Research in Information Assurance and Security

PURDUE UNIVERSITY

A Robust One-Class Bayesian Approach for Masquerade Detection

Qifan Wang and Luo Si
{qifan, lsi}@purdue.edu

Introduction

Masquerade attack is one of the most serious and dangerous security problems, which represents a class of insider or outsider attacks that can occur in many different ways. Detecting masquerade attacks has two main challenges. Firstly, user may not share his/her data with other users for privacy-related reason. Secondly, the training data for each user may be small and results parameter uncertainty. One-class modeling brings certain advantages over two-class modeling for profiling user behavior. One-class modeling requires only user data from the user whose behavior is being profiled, whereas two-class modeling needs data from all other users. In this paper, we propose a robust one-class bayesian approach which takes account of model uncertainty by integrating out the unknown model parameters, while previous methods use a single point estimate to find an optimal model by maximizing the posterior. We also derive the full analytical solution of the predictive distribution over all possible model parameters.

One-Class Bayesian Approach

The problem of small data samples and resulting parameter uncertainty suggests the use of Bayesian techniques. Such an approach offers a natural and principled way to take account of uncertainty by integrating out the unknown model parameters. Although the mode of the posterior could be achieved by existing max-posterior methods, a more powerful approach is to take account of posterior uncertainty when evaluating the probability of a test block *q* belong to the training data *d*, by computing the predictive distribution:

$$P(q|d) = \int_\theta P(q|\theta)P(\theta|d)d\theta = \frac{1}{P(d)}\int_\theta P(q|\theta)P(d|\theta)P(\theta)d\theta$$

In most machine learning techniques, the choice for a prior is the natural conjugate of the likelihood distribution. As in this case, the natural conjugate prior of a *multinomial* distribution is the *Dirichlet* distribution. Under this prior we can compute the resulting posterior, which is also a *Dirichelet* distribution:

$$P(\theta) = Z_\alpha \prod_{i=1}^m (\theta_i)^{\alpha_i - 1} \qquad P(\theta|d) = \frac{\Gamma(|d|+\alpha)}{\sum_{i=1}^m \Gamma(d_i+\alpha_i)} \prod_{i=1}^m (\theta_i)^{d_i+\alpha_i-1}$$

From the posterior distribution, we can derive the predictive distribution as follow equation and then given a test block *q* , by comparing *P(q|self)* with *P(q|non-self)* using the above equation, we can classify *q* as a masquerader's command block when *P(q|non-self*) has larger probability.

$$P(q|d) = Z_q Z_{d+\alpha} \int_\theta \prod_{i=1}^m (\theta_i)^{d_i+\alpha_i-1} d\theta = Z_q \frac{\Gamma(|d|+\alpha)}{\sum_{i=1}^m \Gamma(d_i+\alpha_i)} \frac{\prod_{i=1}^m \Gamma(q_i+d_i+\alpha_i)}{\Gamma(|q|+|d|+\alpha)} = Z_q \frac{\prod_i \prod_{k=1}^{q_i}(d_i+\alpha_i+k-1)}{\prod_{j=1}^{|q|}(|d|+\alpha+j-1)}$$

Experimental Results

Schonlau Data Set (available at ***http://www.schonlau.net***)
Description: 70 users, each contains 15,000 commands. The first 5,000 commands of each user are "clean data"(legitimately issued by the user), and the next 10,000 commands of the target users were randomly injected.
Goal: Detect the masquerade attacks, or "dirty blocks" in the testing commands.

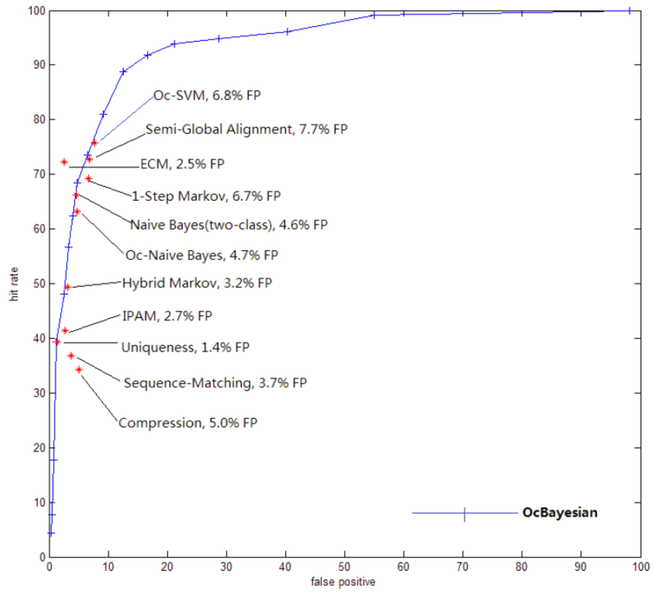

Figure 1: ROC curve for one-class Bayesian model. The best outcomes from other algorithms are also included for comparison.

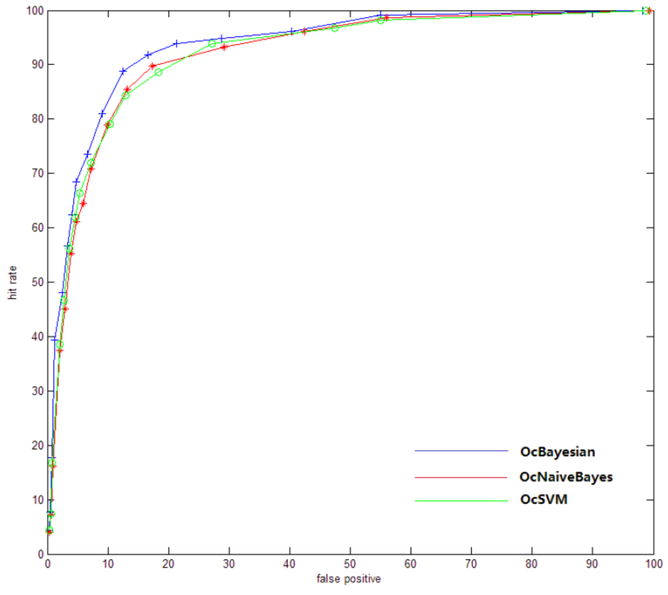

Figure 2: ROC curves of OcBayesian, OcNaiveBayes and OcSVM algorithms.

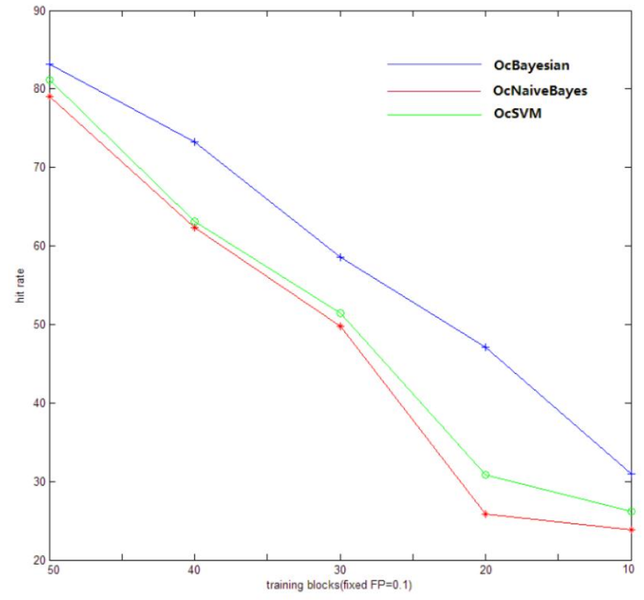

Figure 3: Accuracy of OcBayesian, OcNaiveBayes and OcSVM in different size of training data.

CERIAS®

Discovery Park e-Enterprise Center