



WHITE PAPER

VeriSign[®] Fraud Detection Service

Prepared By:

VeriSign Authentication Services

Product Management

Date: 28 March, 2006



Where it all comes together.™

VeriSign[®] Fraud Detection Service

Prepared By:
VeriSign Authentication Services
Product Management
Date: 28 March, 2006

Identity theft scams have recently become a major and fast-growing method of fraud. The main difficulty in detecting this type of fraud is that the perpetrators operate under the identity of legitimate users, performing apparently legitimate actions. However, although fraudulent intentions cannot be observed directly within the actions, they are reflected by the usage of the stolen identity.

Detection of fraud may be achieved by checking for suspicious changes in a user's behaviour. In practice, we need to observe the user's transactions, discover his or her typical behaviour patterns and take action in the event of an abnormal transaction. The main difficulty is that in most cases there is no specific transaction history for the user that has been labelled as "fraudulent" or "non-fraudulent". Formally speaking, the user's transaction data is unsupervised rather than supervised.

Because of the nature of the data, traditional machine-learning techniques such as decision trees and neural networks, which are applicable to supervised data, are not suitable here. These techniques rely on an early training stage in which input data is labelled, in this case as fraudulent or non-fraudulent. Even if the data happens to be labelled, it will usually be unbalanced in the sense that the number of legitimate actions will be overwhelmingly larger than the number of illegitimate actions. In that case, the illegitimate actions surely do not represent all fraud possibilities. Traditional techniques may detect fraudulent actions similar to ones already recognised as fraud, but they will rarely detect fraudulent activities that were not learned beforehand.

The VeriSign approach for discovering fraudulent behaviour is based on the concept of unsupervised learning. The system itself is required to decide which of the user's actions correspond to his natural behaviour and which are exceptional, without any assistance. With the help of unique clustering algorithms, the VeriSign Fraud Detection System is able to detect suspicious activity within the data in a non-prescriptive way. While the system observes the user's transactions, it discovers common behaviour patterns by grouping similar transactions together. For example, if the user tends to perform his financial activity from work on working days and from home during the weekends, two clusters will be formed, characteristic of this user's common behaviour during both periods:

In order to discover anomalous transactions, new transactions are compared with the user's common behaviour patterns. A transaction that does not correspond with one of these patterns will be treated as a suspicious activity and trigger precautionary steps accordingly.

Unlike other unsupervised statistical techniques used for detecting fraud, our clustering algorithms consider the closeness between similar transaction attribute values, not just between identical ones. Various similarity measurement methods are utilised across the chronological, geographical and financial characteristics for comparing the transactions. These methods reach a higher level of granularity, and thus yield more accurate results. For example, a £1,000 wiring transaction from the City of London at 10:00 AM is perceived as closer to a £1,350 bank transfer from Edinburgh at 10:50 AM than to a £15,000 bank transfer from Bristol at 11:50 PM.

Numerous clustering algorithms have been developed to address different types of problems. Basically, there are two main families of clustering methods: partitional clustering and hierarchical clustering. Partitional clustering methods attempt to divide the given set of elements directly, according to a requested number of clusters. K-means is a typical partitional clustering algorithm. It relies on the ability to measure n elements against k centroids representing k clusters. Each centroid represents the mean (average) of the elements in its cluster. Such methods may yield satisfactory results for data containing numeric attributes, but they are not appropriate for categorical attributes (what is the average of the values “Red”, “White” and “Blue”?).

However, the basic hierarchical clustering methods are not sufficient for accurate results because their grouping of the elements may be imbalanced, forming too large/small or too few/many clusters.

The measurement of similarity between two clusters is performed by one of the following methods: single-link, complete-link and average-link. In single-link measurement the distance between two clusters is considered to be equal to the shortest distance from any element of one cluster to any element of the other cluster. The disadvantage of this measure is that it can generate clusters in which elements are not similar enough due to the fact that it only takes the two most similar elements into consideration. In complete-link measurement, the distance between two clusters is considered to be equal to the longest distance from any member of one cluster to any member of the other cluster. The disadvantage here is that it is bound to generate too many clusters due to our strict unification conditions. In average-link clustering, the distance between two clusters is considered to be equal to the average distance from any member of one cluster to any member of the other cluster. This measure may also generate clusters with dissimilar elements in cases when high similarity averages are accompanied with high variance.

VeriSign uses agglomerative hierarchical clustering algorithms based on the notion of neighbours and links, which overcomes the problems mentioned above. In our implementation, two data elements are considered as neighbours if their similarity upon a domain expert or similarity matrix exceeds a certain threshold. At first, all n data elements are mapped to n clusters respectively. Then, with each iteration, we merge between the two closest clusters such that both clusters fulfil the maximum value of $\text{Link}(C_i, C_j)$, for any pair of clusters C_i and C_j . The metric $\text{Link}(C_i, C_j)$ represents the number of common neighbours between every element in the first cluster to every element in the second one. This measure is normalised by the number of potential neighbours in both clusters, so that a large cluster will not swallow every other cluster and end up with all the elements. Grouping the data elements using links injects global knowledge into the clustering process, forming an optimal division between the elements. Thus, the formed clusters are not too large or too small, and they contain elements which are relatively similar one to another.



Since we cannot predict the number of user behaviour patterns in advance, our algorithms were developed to produce the real number of clusters, representing each of the user's behavioural patterns, as they actually appear within the data. In addition, the maximum number of clusters can be defined so that if similarity thresholds were strictly set, the algorithms would reduce them in order to generate a reasonable number of clusters.

Similarly, we introduce the notion of clustering execution levels, which lets us define different clustering configurations for various situations of available data. If, for example, we start the anomaly detection process with a relatively sparse data set, we may wish to activate the clustering phase with reduced similarity thresholds, since the number of common neighbours between pairs of data elements is bound to be small. This enables us to perform fraud detection at early stages of the data collection, without generating false positives or false negatives. In each execution level, we define the minimum number of transactions, similarity threshold, confidence factor and the participating attributes.

Visit us at www.verisign.co.uk. for more information.

© 2006 VeriSign UK Ltd. All rights reserved. VeriSign, the VeriSign logo, "Where it all comes together", and other trademarks, service marks and designs are registered or unregistered trademarks of VeriSign and its subsidiaries in the United States and in foreign countries. All other trademarks are the properties of their respective owners.

00023299 28/03/06
ML 061066