FOREWARD

These working notes contain the technical papers presented at the *Workshop on AI Approaches to Fraud Detection and Risk Management*, held July 27, 1997, in conjunction with the 1997 National Conference on Artificial Intelligence (AAAI-97).

The purpose of the workshop was to gather together researchers and practitioners in the areas of risk management, fraud detection, and network/computer intrusion detection, in order to discuss common issues in the application of intelligent technologies to these problems, to share approaches and techniques, and generally to come to a better understanding of both the problems and their solutions. Our intent was to facilitate dialog between related application areas, as well as between researchers and practitioners. To our knowledge, this workshop was the first forum bringing together researchers and practitioners working in these three related areas.

Risk management, fraud detection, and intrusion detection all involve monitoring the behavior of populations of users in order to estimate, detect or avoid undesirable behavior. This workshop focused specifically on improper behavior, which includes fraudulent usage, intrusion, delinquency, and account defaulting.

Fraud detection and risk management problems share characteristics that complicate the application of existing AI technologies. Solving such problems involves processing very large volumes of data, more than most AI techniques are designed for. The proportions of legitimate and undesirable behavior may be very skewed, sometimes by orders of magnitude, and the proportions may change over time. In addition, decisions made by such systems must be sensitive to error costs, which may change over time as well. Both legitimate and undesirable behavior tend to drift over time, and any static system will be soon be obsolete. Finally, researchers in these domains must also be sensitive to social issues such as privacy and discrimination.

The working notes contain sixteen papers, ten of which were selected for presentation at the workshop. As reflected by the presentation schedule, these ten papers were grouped into three categories.

- Four papers addressed issues in applying classification techniques to fraud and risk problems. The issues include: the use of clustering techniques to generate class labels, the use of techniques from decision analysis and ROC analysis to deal with uncertain costs and class distributions, and the use of meta-learning techniques to enable information hiding.

- Three papers presented approaches to modeling legitimate behavior for the detection of anomalous activity in order to detect fraud or intrusions. These papers addressed the problem that examples of improper behavior may be scarce, and that any database of improper behavior is unlikely to be complete.

- Finally, three papers looked beyond currently implemented systems for a view of what future systems should address. In particular, systems should be able to deal with data at many different levels of aggregation (transactions, sequences of transactions, accounts); they should look beyond these to groups of accounts and transactions, for example to identify collusive agents; they should be able to transcend traditional systems-oriented boundaries, and they may be better viewed as investigative tools, instead of stand-alone solutions.

Our hope was for the workshop to facilitate interaction between researchers and practitioners, focus on the commonalities among fields previously treated in isolation, and cross-fertilization among the fields. We thank the authors and attendees for their efforts and enthusiasm in making this possible. We are indebted to the AAAI organization for organizational and funding assistance, and for publishing these working notes; to Ray Mooney, Chair of the AAAI-97 Workshop Committee; to Mary Beth Jensen, AAAI workshop co-ordinator; and to our anonymous workshop-proposal reviewers for their suggestions and encouragement.

Tom Fawcett, Ira Haimowitz, Foster Provost and Salvatore Stolfo