Prospective Assessment of AI Technologies for Fraud Detection: A Case Study

David Jensen
Computer Science Department
Campus Box 34610 LGRC
University of Massachusetts
Amherst, MA 01003-4610
jensen@cs.umass.edu

Abstract

In September 1995, the Congressional Office of Technology Assessment completed a study of the potential for AI technologies to detect money laundering by screening wire transfers. The study, conducted at the request of the Senate Permanent Subcommittee on Investigations, evaluates the technical and public policy implications of widespread use of AI technologies by the Federal government for fraud detection. Its conclusions are relevant to many other uses of AI technologies for fraud detection in both the public and private sectors.

Keywords: wire transfers, money laundering, evaluation, technology assessment.

Introduction

In January of 1994, the Senate Permanent Subcommittee on Investigations requested a study of the feasibility of using AI technologies to detect money laundering in international wire transfers. The Subcommittee made the request to the Office of Technology Assessment (OTA) — an analytical support agency charged with producing in-depth studies of science and technology issues for Congress. The author was one of three OTA analysts who conducted the study and wrote the final report.[1]

The study and its findings hold lessons for researchers, developers, and managers of AI-based systems for detecting financial fraud. The study is one of the earliest broad-scale assessments of the technical potential and policy implications of using AI technologies to detect financial fraud. While the study analyzes a specific task and set of technical configurations, many of its conclusions apply to other systems. For example, it identifies a number of technical challenges with broad relevance, including high data volume, high false positive rates, changing profiles, adaptable adversaries, and the similarity of legitimate and fraudulent behavior.

This paper presents an overview of the OTA study, with special focus on the technical challenges it identifies. The first section discusses the basics of the study, the second section introduces money laundering, the third section introduces wire transfers, the fourth section discusses the technologies reviewed by the study, and the final section outlines some of the study's findings and the personal conclusions of the author.

The Study

The Subcommittee's interest in the study arose from several sources. First, some evidence suggested that criminal organizations were making use of the international banking system to transfer their profits overseas, beyond the reach of U.S. law enforcement agencies. Second, improved technologies for semi-automated data analysis and screening were being deployed by the U.S. intelligence community, other government agencies, and private industry. The possibility existed that such technologies could be applied to screen large volumes of financial transactions for evidence of money laundering. Finally, there was mounting evidence that several related efforts were already successful at identifying financial crimes. An agency of the U.S. Department of Treasury — the Financial Crimes Enforcement Network (FinCEN) — already screened reports of large currency transactions at banks and retailers, called Currency Transaction Reports (CTRs) (Senator *et al.* 1995; Goldberg & Senator 1995). A similar agency in Australia — the Australian Transaction Reports and Analysis Centre — was successfully screening wire transfer records.

Armed with the Subcommittee's request, OTA began its study in January of 1994. OTA's approach to

[1] U.S. Congress, Office of Technology Assessment, *Information Technologies for the Control of Money Laundering*, OTA-ITC-630 (Washington, DC: U.S. Government Printing Office, September 1995). The full text of the report is available in paper from the Government Printing Office and electronically from http://www.wws.princeton.edu/~ota/.

conducting a study was relatively standard: A small group of analysts conducted interviews, held workshops, visited relevant government and private offices, and read extensively about the relevant subject matter. In parallel, the agency formed an advisory panel of representatives from relevant communities. The advisory panel neither wrote nor approved the report, but offered comments, recommended contacts, and suggested basic directions for the analysis. Finally, the team wrote a draft report, sought extensive review from the advisory panel and others, revised the draft, and issued a final report. In all its studies, OTA's role was not to recommend specific policy actions to the elected representatives in Congress. Rather, OTA identified relevant facts that support the creation of public policy and identified areas where there is still widespread disagreement.

This study was fairly typical. The project team consisted of three analysts with expertise in, respectively, technology assessment and financial computer systems, law, and artificial intelligence. The study was completed in approximately 18 months and involved scores of interviews, numerous site visits in Washington and New York, and workshops on wire transfers, privacy and confidentiality, AI technologies, and possible system configurations. The advisory panel drew members from the banking, law enforcement, legal, and technical communities. Reviewers and contributors to the report came from government (including the U.S. Departments of Treasury and Justice, Federal and state law enforcement agencies, White House Office of Narcotics and Drug Control Policy, and The Federal Reserve), financial institutions, universities, and public interest groups.

Money Laundering

A central issue of the study was the relationship between a criminal activity — money laundering — and a widely-used businesses tool — international wire transfers.

To launder money is to disguise the origin or ownership of illegally gained funds to make them appear legitimate. Hiding legitimately acquired money to avoid taxation also qualifies as money laundering. Federal agencies estimate that as much as $300 billion are laundered annually. Of this, some $100 billion are thought to be drug profits, and $40 billion to $80 billion of that are generated in the United States. These estimates are approximate at best, representing a mix of experience, intuition, and extrapolation.

Money laundering is required because many crimes generate large cash profits. These profits can raise suspicions of law enforcement agents and lead to arrest if they are not made to appear legitimate. The money itself is also subject to seizure by local, state, and federal authorities, depriving criminal organizations of their profits even if most of the organization's members evade arrest.

One of the goals of money laundering is to convert cash proceeds to another form. Although some illegal proceeds are used locally, many criminal organizations transfer profits outside the country. This puts the money beyond the easy reach of U.S. law enforcement. In the case of drug proceeds, it also returns profits to the original producers, who are often located outside the United States. Cash is bulky and relatively easy to detect, so it is not the preferred method for international transfer. Instead, money launderers attempt to convert cash to more easily transferred forms.

Law enforcement officials often describe three stages of money laundering:

- Placement — introducing cash into the banking system or into legitimate commerce;
- Layering — separating the money from its criminal origins by passing it through several financial transactions;
- Integration — aggregating the funds with legitimately obtained money or providing a plausible explanation for its ownership.

U.S. Department of Treasury had already attacked the first stage — placement. Financial institutions and merchants are required to file Currency Transaction Reports (CTRs) when certain types of large currency transactions are made. For example, CTRs are generated when a customer deposits more than $10,000 in cash or purchase a car with more that $10,000 in cash.

However, measures to prevent placement have not been entirely successful, and some evidence suggested that wire transfers are used in the second stage — layering. Because of this potential "silent pipeline" for transferring illegal proceeds out of the country, interest arose within Congress and the Executive branch in identifying wire transfers used for money laundering.

Wire Transfers

Wire transfers are the primary mechanism used by the business community for fast and reliable transfer of funds between two parties. The simplest funds transfers occur between banks that maintain a "correspondent" relationship — banks that each maintain an account in the name of the other bank. In this case, a secure message between the banks can result in a "book transfer" where funds are simultaneously debited from one account and credited to another. Two

banks that do not have a correspondent relationship can still transfer funds if they can establish a chain of banks that do have such a relationship. This latter process is eased by the existence of 15 to 20 large "money center" banks that maintain correspondent relationships with smaller banks and with each other. This process of book transfers is directly analogous to wire transfers.

Two systems can be used to send and receive wire transfers: Fedwire, operated by the Federal Reserve Banks; and CHIPS, the Clearing House for Interbank Payments System.[2] In 1994, Fedwire and CHIPS processed nearly 120 million transactions with a total dollar value of more than $500 trillion dollars. Although those transactions amount to only about 0.1 percent of all payments in the United States (the vast majority of payments are by check), they were responsible for more than 90 percent of the dollar value of all payments.

Wire transfers are rarely used by individuals. They are primarily used by large corporations sending large-dollar transfers. Most of these transactions are handled automatically and only seen by human operators if they contain formatting errors. Legitimate businesses use wire transfers when sending very large sums of money, and when timeliness and certainty are of paramount importance.

Technologies and Technical Issues

The study examined technologies for three activities: knowledge acquisition, knowledge use, and data exploration. Technologies for *knowledge acquisition* draw heavily from the fields of machine learning, knowledge discovery, and statistical modeling. Several proposals for using AI technologies for wire transfer screening involved the use of knowledge acquisition techniques to develop profiles of money laundering. While law enforcement agencies are sometimes quite knowledgable about money laundering in general, they lack reliable profiles that can distinguish illicit wire transfers from legitimate ones. Knowledge acquisition technologies were proposed by advocates as a way to produce the requisite profiles.

Technologies for *knowledge use* include knowledge-based systems and techniques for knowledge management. Once reliable profiles of illicit wire transfers were developed, many proposals called for deployment of those profiles in knowledge-based systems. The proposals deployed these systems in banks, wire transfer systems, or law enforcement agencies. In addition, knowledge management techniques were explored because the necessary profiles are expected to change frequently, as money launderers adapt to new detection capabilities. This necessitated technologies for revising, extending, and deploying profiles to a potentially large number of locations in the United States (e.g., all banks that make wire transfers).

Technologies for *data exploration* include visualization and link analysis. The perceived importance of these technologies increased throughout the study, as the team visited and interviewed the staff of existing organizations that investigate financial crimes. Nearly the only commonality among the Federal and state law enforcement agencies investigating financial crimes was the use of a single analytical technique: link analysis. Link analysis examines a large number of potentially dissimilar database records and establishes links among those records based on data fields with identical values (Sparrow 1991). For example, a record of a person might be linked to a record of a business based on a match between the person's work address and the address of the business. Similarly, several persons might be linked based on a shared bank account or because they live on the same street. Link analysis is an indispensable tool to visualize and evaluate networks of people, places, and things.

Examining these three classes of technologies raised several important issues:

- *High data volume:* CHIPS and Fedwire handle approximately 120 million transactions per year, or about 500,000 transaction each business day. This volume of transactions is far larger than the 30,000 transactions per day currently provided to the Financial Crimes Enforcement Network (FinCEN), the Treasury agency responsible for analyzing Currency Transaction Reports (CTRs).
- *Low incidence of illegitimate transactions:* The number of wire transfers involving money laundering is quite small compared to the overall volume of wire transfers. Even under generous assumptions, money laundering would account for less than one-tenth of one percent of all wire transfers. This low incidence could lead to an extremely high false positive rate in a deployed system. If the system is 95% accurate, and the errors are distributed proportionally, then 99% of the transfers identified as illicit would actually be legitimate.
- *Lack of tested profiles:* Law enforcement agents, bankers, and bank regulators readily admit that they cannot supply profiles that reliably distinguish between legitimate and illicit wire transfers.

[2] A third system, SWIFT, is sometimes added to this list, although it is not an electronic funds transfer system, but rather a specialized international cooperative communications system. However, most SWIFT messages of importance for this paper result in a corresponding CHIPS message. See the full study for details.

- *Lack of labeled cases:* It is difficult to label individual wire transfers as definitely associated with money laundering. Years often elapse between initial investigation and final prosecution, and it is unlikely that such analysis would detect all, or even most, of the transfers associated with money laundering. This differs from some other forms of fraud that are "self-revealing." For example, some cellular phone fraud becomes evident shortly after it is perpetrated because customers complain about unexplained charges on their bills. This lack of labeled cases severely limits the applicability of many techniques for machine learning and knowledge discovery.

- *Intelligent adversaries:* Criminal organizations are highly adaptive, and there are many ways to launder money. One convicted money launderer interviewed by OTA insisted that criminal organizations "instantly" know when money laundering detection methods are changed, because they have informants within banking, law enforcement, and intelligence communities. Any system using static profiles will quickly become outdated as money launderers change their tactics in response to detection systems.

Findings and Observations

After reviewing the technical, economic, social, legal, and international issues surrounding money laundering and wire transfers, OTA reached the following conclusions:

- OTA was unable to conceptualize any system without substantial social and economic costs. However, given the importance of controlling international crime, some social and economic costs may be acceptable.

- The simplest system configuration — continual, automated, real-time screening of all wire transfers — would probably not be effective in detecting money laundering. Identifying one illicit transaction among thousands of legitimate transfers is difficult or impossible based on information from the transfers alone.

- Two more complex system configurations may be feasible. One would allow a designated federal agency to electronically request access to specific wire transfer records stored at banks. The other involves banks automatically forwarding a small subset of wire transfer records to a designated federal agency for further analysis. Both involve moderate to high costs and raise serious privacy issues.

- The development of new forms of electronic payments — often referred to as "digital cash" — threatens to undercut money laundering controls based on screening currency transactions and wire transfers. However, the impact of digital money on money laundering is still highly uncertain.

Additional detail on these and other conclusions can be found in the published report.

In addition to helping formulate the findings of the study, the author reached a set of personal conclusions. These do not appear in the published report, because they concern methodology and system-building. While neither type of conclusion belonged in the OTA report, they may be useful to system designers and managers. The conclusions are:

- Prospective evaluation of fraud detection systems can be difficult, but informative. Such evaluation should consider the utility and design of analogous systems, the available data and knowledge about the task, the expected output, and the potential of human investigators to use that output.

- The existence of necessary knowledge and information is one of the most important factors to evaluate prospectively. Many AI systems require either profiles of fraudulent activity or labeled data to develop such profiles. Lacking explicit profiles, or the data to develop them, implies a radically different system design.

- Designers and users of existing analogous systems are some of the best sources of information. For example, OTA interviewed designers and users of systems for law enforcement analysis, commercial fraud detection, and intelligence. In many cases, their knowledge of system design was implicit, and had to be derived from interviews and demonstrations, rather than simple requests. However, it was among the most important information the team uncovered.

- While the term *fraud detection* is often used to describe a desired system, most deployed systems have evolved into systems for *fraud investigation.* This evolution is often accompanied by a shift from an automated to an interactive system. The latter type of system can still use AI tools, but deploys those tools to aid human investigators.

- Many of the most useful systems for fraud investigation involve more than just the ability to store and analyze individual records. For example, nearly every major system for detecting financial fraud provides tools to construct and analyze networks of related records (e.g., all records that share a common

account number or common address). Other systems provide the ability for analysts to annotate records and preserve those annotations to aid future analyses.

Acknowledgments

The author wishes to thank his colleagues in conducting the study, Vary Coates and Steven Bonorris. The opinions expressed here do not represent those of the Office of Technology Assessment or the United States Congress.

References

Goldberg, H. G., and Senator, T. E. 1995. Restructuring databases for knowledge discovery by consolidation and link formation. In *Proceedings of the First International Conference on Knowledge Discovery and Data Mining (KDD-95)*, 136–141. AAAI Press.

Senator, T. E.; Goldberg, H. G.; Wooton, J.; Cottini, M. A.; Khan, A. U.; Klinger, C. D.; Llamas, W. M.; Marrone, M. P.; and Wong, R. W. 1995. The FinCEN Artificial Intelligence System: Identifying potential money laundering from reports of large cash transactions. In *Innovative Applications of Artificial Intelligence 7*, 156–170. AAAI Press.

Sparrow, M. K. 1991. The application of network analysis to criminal intelligence: An assessment of the prospects. *Social Networks* 13:251–274.