RIVISTA ITALIANA DI INFORMATICA E DIRITTO
PERIODICO INTERNAZIONALE DEL CNR-IGSG

ISSN 2704-7318
Fascicolo 2-2022
DOI: 10.32091/RIID0092
Articolo sottoposto a peer review

Efficient and trustworthy decision making through human-in-the-loop visual analytics: A case study on tax risk assessment

Walter Didimo • Luca Grilli • Giuseppe Liotta • Fabrizio Montecchiani

Data mining and AI techniques are increasingly being used to automate data analysis. Ideally, one may wish to completely automate the data analysis process, but in many real-world applications a full automation may pose significant risks. In these cases, human analysts must be directly involved to refine the analysis or to make the final decisions. A challenging problem, therefore, is how to perform efficient and trustworthy decision-making when humans are an integral part of the analysis pipeline. We propose a "human-in-the-loop" methodology that leverages data mining, machine learning, and visual analytics to improve and speed up the analysis. A key feature is the use of a dashboard that integrates intuitive visual tools, which aid analysts to efficiently discover hidden data patterns or to get helpful insights. We describe in particular how this methodology has been successfully applied to support Revenue Agency officers in tax risk assessment.

Human-in-the-loop – Visual analytics – Data mining – Tax risk assessment

1. Introduction

In a digitized world, the extraction of knowledge from data is a fundamental need to support decisions or predict the behavior of observed phenomena. Advances in data mining, machine learning and AI play a critical role in optimizing this process, and potentially might completely automate it, including the decision-making phase. Figure 1 shows an illustration of a fully automated data analysis. Data from various sources are collected and stored based on appropriate data models. Specific data portions, in a predefined format, are then automatically extracted and passed to the data processing phase, whose outcome is a prediction or a decision that triggers some action. No human is involved in this analysis pipeline, apart from its design, configuration and periodic supervision. Today, there are several software applications or features that work this way. For example, anti-spam filters automatically classify an incoming email as spam, and move it to the junk folder, without a direct interaction with the user.

Attempts to completely automate data-driven decision-making are being made in several fields such as finance, e-commerce, health, automotive, cybersecurity and the environment. For instance, many cybersecurity companies are investing a lot of resources to develop security incident response automation[1]. They wish to provide systems that can respond to cyberattacks at computer speeds – with a very limited human involvement – by exploiting machine learning and AI techniques.

The Authors belong to the Department of Engineering, University of Perugia. This work was supported in part by the University of Perugia, Fondi di Ricerca di Ateneo 2021, Project "AIDMIX—Artificial Intelligence for Decision Making: Methods for Interpretability and eXplainability". This essay is part of the special section on *Gli algoritmi pubblici tra legalità e partecipazione* edited by Benedetto Ponti.

However, at least in the short term, a full automation may pose significant or unacceptable risks, in many application domains. Apart from (possible) ethical problems, the main issue is about the level of uncertainty that often characterizes the data in input to the processing phase, which, in turn, propagates to its outcome. In other words, the uncertainty may drastically compromise the accuracy of automated analysis. It is also unclear how machine learning algorithms may react to inputs that significantly differ from those used in the training phase. Because of this uncertainty, an AI-based software that automates clinical diagnoses can pose serious risks to patients' health and safety, if no rigorous human validation of its outcomes is carried out. In these cases, the data analysis process must necessarily include the presence of a human analyst, who is the responsible for making the final decisions or for refining the analysis[2].

There exists an extensive body of literature devoted to describe the role that humans play in the development of effective machine learning algorithms, which is referred to as *human-in-the-loop machine learning* (*HILML*)[3].

More in general, human-in-the-loop approaches are not only restricted to machine learning, but they are attracting increasing interest in data science. Here, the idea is to make humans an integral part of the data analysis pipeline, since some tasks may require the sensitivity of a human analyst with high experience in the specific application domain. This emerging topic, known as *human-in-the-loop data analysis* (*HILDA*)[4], focuses on how to optimize the human-computer cooperation through the usage of suitable user interfaces and interaction models. In particular, when using ML or AI engines in data analysis systems, human analysts may strongly benefit from visual analytics techniques to understand the decisions or predictions made by these engines[5].

Information visualization tools are therefore key components in the fields of HILDA and explainable AI, since they provide a valuable aid in the design of decision support systems based on the human-in-the-loop paradigm. According to this paradigm, the user constantly interacts with the system with the aim to: (*i*) convey to the system their knowledge and experience in order to provide rules or tuning mechanisms for the machine learning engines; (*ii*) inspect the results automatically provided by the system in order to understand the nature of such results and validate their appropriateness, reliability, and usability, depending on the specific domain of interest.

Hereafter, we propose a human-in-the-loop methodology that leverages visual analytics techniques to improve and speed up the data analysis process. The methodology stems from the experience gained in the development of three decision support systems for tax risk assessment, whose effectiveness has been experimentally validated in previous works[6].

2. Human-in-the-loop methodology for data analysis

Our human-in-the-loop methodology is illustrated in Figure 2 and relies on the following principles:

- *final human validation*: final decisions are always made by a human analyst, machine learning and AI modules can be used, but the analysis cannot be completely automated, and a human mediation is mandatory;
- *human-centered analysis*: the analyst plays a central role in the data analysis process, providing valuable knowledge and expertise, in particular, they interact with the system throughout the process by:
  - (i) selecting the data sources;
  - (ii) defining the analysis entry point;
  - (iii) choosing the data processing algorithms;
  - (iv) exploring the results of the system;
  - (v) validating and refining the results.
- *visual analytics-driven workflow*: the analyst interacts with a UI dashboard, in a domain-specific software system, to carry out every analysis task; the dashboard integrates sophisticated visualization techniques and visual tools to speed up the analysis, to discover hidden data patterns or to get helpful insights.

The proposed methodology gives a very abstract outline of how the analysis process should be conducted. In particular, the goals and the specific analysis tasks, as well as the type of data itself may drastically impact on the UI, on the visual analytics techniques, on the data model, and on the processing algorithms. Therefore, for a more concrete understanding, we will describe a case study of our methodology to support Revenue Agency officers in the tax risk assessment[7].

3. Case study: tax risk assessment

Revenue Agency (RA) officers devote a lot of work to discover tax evasion and tax avoidance. The *fiscal audit process* refers to the set of activities they carry out for this purpose. This process largely consists of the analysis of data within the RA database, which includes information about single taxpayers and their relationships like *economic transactions*, *shareholdings*, and *corporate positions*. Specifically, analysts look for certain interactions among taxpayers, called *suspicious patterns*, that could reveal tax evasion activities. Suspicious patterns are derived by abstracting the characteristic traits of real cases of tax evasion/avoidance or, more in general, of fiscal risky schemes. They can be conveniently modeled as topological structures, i.e. as small groups of interconnected entities, along with the specification of some constraints on economic attributes associated with these entities and/or with their interconnections. For example, Figure 3a shows a real case of a risky scheme involving a company (C1) in the building sector and a real estate company (C2). C1 sells a building to C2, issues an invoice for 1 million euro, declares its VAT debt for this invoice, but never pays it. While C2 requests a refund for the VAT credit of its purchase. By analyzing the refund request, RA officers discovered that C1 and C2 were owned by the same taxpayer (T1). The VAT refund was therefore denied by the RA, due to the high fiscal risk exhibited by this scheme. A specific suspicious pattern, named SuppliesFromAssociated, has been derived from this real case of risky scheme (see Figure 3b for an illustration). This pattern has a triangular topology, where both the seller (C1) and the buyer (C2) of an economic transaction are participated by the same subject (T1). To generalize the pattern, the following additional rules have been included: the amount of the economic transaction is greater than or equal to a specified threshold Y; the VAT payment of the seller is less than a specified threshold Z; the VAT credit declared by the buyer is higher than a specified threshold X.

Since suspicious patterns have an inherent topological structure, their detection can be significantly improved by adopting a *taxpayer network* data model, which can be built from the RA (relational) database and stored into a graph database[8]. In a *taxpayer network* $\mathcal{N}$, each node ν of $\mathcal{N}$ represents a single taxpayer, i.e. either an individual or a legal person, like a private company or a public institution. Whereas arcs in $\mathcal{N}$ are directed and model economic relationships between pairs of taxpayers. For example, an arc (μ, ν) may represent an economic transaction where taxpayer μ is the seller and taxpayer ν is the buyer.

Ideally, a fully automated fiscal audit could work as follows. Let E be the total set of actual tax evasion attempts or events that can be revealed by an optimal analysis of the taxpayer network.

– A first set $C \supset E$ of candidate risky schemes is automatically extracted from $\mathcal{N}$.
– Then, every element in $C \setminus E$ is (automatically) filtered out by performing a suitable data processing on the input C.
– Eventually, the final output is an automated generated report describing each element of E.

Unfortunately, due to various sources of uncertainty, the aforementioned automated analysis fails, namely it may produce many false positives and false negatives. The main sources of uncertainty are listed below.

– *Missing or imprecise data sources*: a taxpayer network typically contains several hundred thousand nodes and some million arcs, consequently some data items could be missing, wrong, imprecise, or outdated.
– *Inherent uncertainty of risky schemes*: the identification of risky schemes necessarily requires an implicit or explicit knowledge of their models (i.e., suspicious patterns), but every model is an approximation of the real world.
– *High dimensionality of $\mathcal{N}$*: there are various kind of nodes and arcs in $\mathcal{N}$, which may have several attributes of different types, therefore the identification of reliable models of risky schemes may require an infeasible number of observations.
– *Unpredictability of future risky schemes*: future risky schemes can be substantially different from those already known, but detection strategies are based on the latter ones.

We therefore propose a human-in-the-loop methodology for discovering actual risky schemes in a taxpayer network. Our methodology focus on how to support human analysts – through a suitable visual analytics system – to efficiently find a set E' of risky schemes such that E' is hopefully a good approximation of E. The final decision about whether a specific risky scheme in E' is an actual tax evasion activity can be made only by human analysts after a rigorous examination of the related documentation; we omit the description of this phase because it does not involve any digital tool. Set E' can be found through several analysis loops that require a strong user involvement. Salient steps of a generic iteration are listed below:

– Suspicious pattern definition
– Suspicious pattern matching
– Automatic preliminary taxpayers' risk assessment
– Visual exploration of matching results.

We now briefly describe each of them in a separated subsection.

3.1. Suspicious pattern definition

Due to the enormous size of a typical taxpayer network, it is not possible to explore it as a whole. Indeed, risky schemes of E' can be found by exploring only small subnetworks $\mathcal{S}_i$ $(1 \leq i \leq |E'|)$ of $\mathcal{N}$ containing taxpayers' interactions that conform to some suspicious pattern. The specification of a suspicious pattern is therefore the entry point for the extraction of these subnetworks of $\mathcal{N}$. Analysts, based on their domain knowledge, can define new suspicious patterns using a high level domain-specific visual language (see, e.g., Figure 4)[9], which significantly speeds up this step. Moreover, a library of known suspicious patterns can be easily created and reused.

3.2. Suspicious pattern matching

Let $E'(P) \subset E'$ be the set of risky schemes associated with a specific suspicious pattern P. The extraction of the corresponding subnetworks $\mathcal{S}_j$ $(1 \leq j \leq |E'(P)|)$ is carried out by running a graph pattern matching algorithm on the whole taxpayer network (see Figure 5 for an illustration). Basically, this algorithm takes in input P and $\mathcal{N}$, and returns all the subnetworks in $\mathcal{N}$ that match P. The use of a graph database to store the taxpayer network makes it possible to efficiently identify subnetworks S_j.

3.3. Automatic preliminary taxpayers' risk assessment

Analysts have to assess the fiscal risk of taxpayers in the subnetworks returned by the previous matching algorithm. This phase can be partially automated by assigning to each taxpayer certain risk indexes, which provide a first estimation of their fiscal risk. Furthermore, risk indexes make it possible to rank taxpayers from the most risky to the least risky. Analysts can therefore start their in-depth investigations from the riskiest ones. Various types of risk indexes are computed, including both classical social network analysis (SNA) indexes and some domain-specific indexes. We also make use of a tax risk forecasting model[10] based on machine learning algorithms. The forecasting model is trained on the basis of the outcome of previous fiscal audits. It turns out to be quite effective on identifying the most risky taxpayers. Furthermore, based on the consensus that risky subjects may negatively affect the behavior of their business partners, we apply an information diffusion method to propagate the fiscal risk in the taxpayer network. The diffusion process is based on a stochastic model that simulates the spread of an information over an underlying network.

3.4. Visual exploration of matching results

Starting from the subnetworks returned by the graph pattern matching phase, analysts carry out a visual exploration of its neighborhood in the taxpayer network, enriched with the fiscal risk scores (see, e.g., Figure 6)[11].

The purpose of this phase is to support the analyst in validating the fiscal risk scores assigned by the previous phase. This human validation activity is fundamental for the tax administration, which must have complete control over the taxpayer selection process. In fact, thanks to a visual exploration of the taxpayer network, the analyst can better assess the real risk profile of taxpayers, thus carrying out a more effective selection of tax audits. Namely, the analyst can find new risky graph patterns or false negative cases. This information closes the loop of the system by enriching the pattern library and by improving the performance of the forecasting model.

4. Final remarks

We have described a human-in-the-loop methodology to support an efficient and trustworthy decision-making process. As clearly emerged in the description of the tax risk assessment case study, an effective application of this methodology requires a strong involvement of domain experts, whose contribution embraces both the usage of decision support systems and the design and tuning of their software modules.

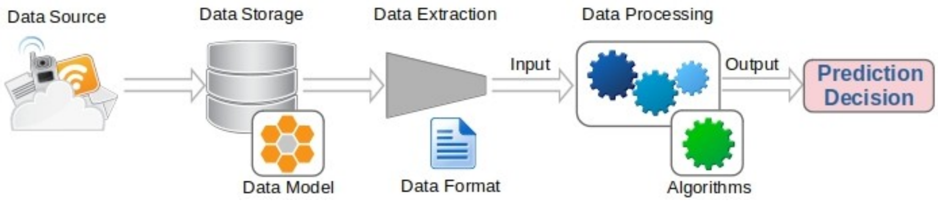

Figure 1: Illustration of a fully automated data analysis

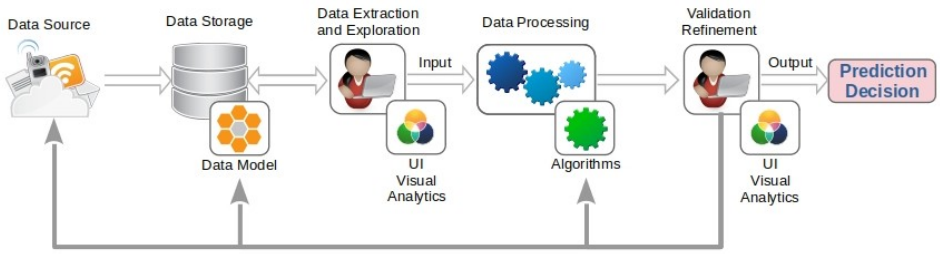

Figure 2: Illustration of a human-in-the-loop methodology for data analysis

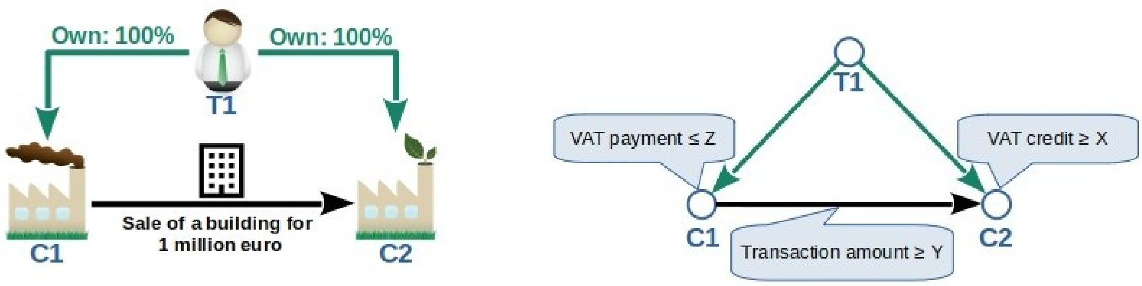

(a) real case of a risky scheme

(b) suspicious pattern SuppliesFromAssociated

Figure 3: Illustration of (a) real case of a risky scheme and of (b) the corresponding suspicious pattern

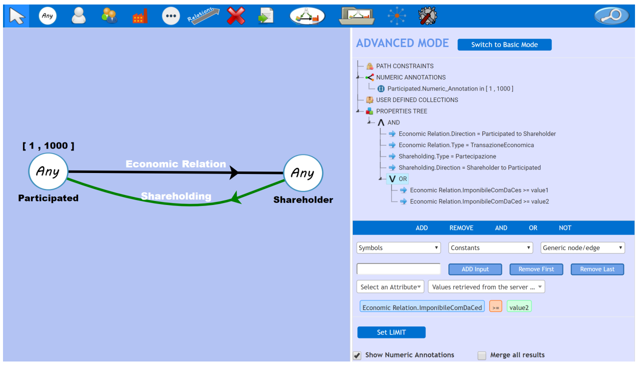

Figure 4: Graphical interface used in TaxNet to define a suspicious pattern

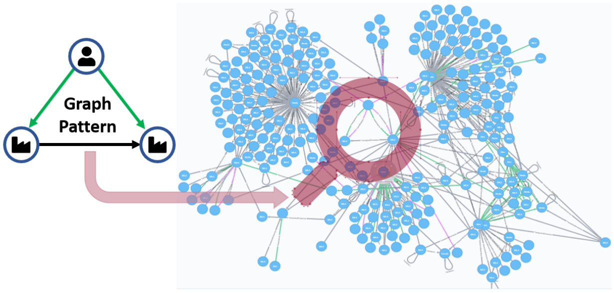

Figure 5: Tax officers encode risky relational schemes among taxpayers into suspicious graph patterns to be searched in the taxpayer network. The pattern in the figure represents a SuppliesFromAssociated scheme, consisting of an economic transaction (black edge) and two shareholding relationships (green edges)

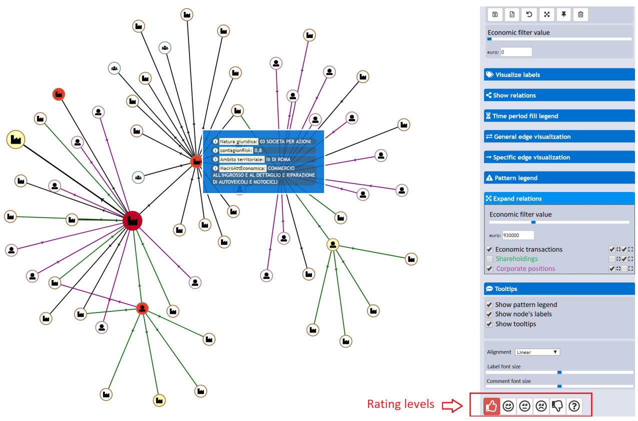

Figure 6: Visualization of a portion of the taxpayer network in MALDIVE Starting from a *high-risk* subject (dark red background), the analyst expands the network by exploding the relationships of some taxpayers. Black edges represent economic transactions, green edges are shareholdings, and purple edges are corporate positions. At the end of the analysis the analyst may add a personal feedback to the record of the investigated taxpayers

Note

[1]C. Lawson, *Market Guide for Security Orchestration, Automation and Response Solutions*, Gartner, 2022.

[2]S. Budd, E.C. Robinson, B. Kainz, *A survey on active learning and human-in-the-loop deep learning for medical image analysis*, in "Medical Image Analysis", vol. 71, 2021, 16 p.

[3]*Ibidem*; R.M. Monarch, *Human-in-the-Loop Machine Learning: Active learning and annotation for human-centered AI*, Manning, 2021; E. Mosqueira-Rey, E. Hernández-Pereira, D. Alonso-Ríos et al., *Human-in-the-loop machine learning: a state of the art*, in "Artificial Intelligence Review", 17 August 2022; X. Wu, L. Xiao, Y. Sun et al., *A survey of human-in-the-loop for machine learning*, in "Future Generation Computer Systems", vol. 135, 2022, pp. 364-381.

[4]A. Doan, *Human-in-the-Loop Data Analysis: A Personal Perspective*, in "Proceedings of the Workshop on Human-In-the-Loop Data Analytics", Association for Computing Machinery, 2018, pp. 1-6.

[5]F. Cheng, Y. Ming, H. Qu, *DECE: Decision Explorer with Counterfactual Explanations for Machine Learning Models*, in "IEEE Transactions on Visualization and Computer Graphics", vol. 27, 2021, n. 2, pp. 1438-1447; J. Choo, S. Liu, *Visual Analytics for Explainable Deep Learning*, in "IEEE Computer Graphics and Applications", vol. 38, 2018, n. 4, pp. 1558-1756; Z. Jin, Y. Wang, Q. Wang et al., *GNNLens: A Visual Analytics Approach for Prediction Error Diagnosis of Graph Neural Networks*, in "IEEE Transactions on Visualization and Computer Graphics", 4 February 2022; H. Li, M. Xu, Y. Wang et al., *A Visual Analytics Approach to Facilitate the Proctoring of Online Exams*, in P. Bjørn, S. Drucker (eds.), "Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems", Association for Computing Machinery, 2021; J. Yuan, C. Chen, W. Yang et al., *A survey of visual analytics techniques for machine learning*, in "Computational Visual Media", vol. 7, 2021, pp. 3-36.

[6]W. Didimo, L. Giamminonni, G. Liotta et al., *A visual analytics system to support tax evasion discovery*, in "Decision Support Systems", vol. 110, 2018, pp. 71-83; W. Didimo, L. Grilli, G. Liotta et al., *Combining Network Visualization and Data Mining for Tax Risk Assessment*, in "IEEE Access", vol. 8, 2020, pp. 16073-16086; Id., *Visual querying and analysis of temporal fiscal networks*, in "Information Sciences", vol. 505, 2019, pp. 406-421.

[7]Further details can be found in: W. Didimo, L. Giamminonni, G. Liotta et al., *A visual analytics system*, cit.; W. Didimo, L. Grilli, G. Liotta et al., *Combining Network Visualization*, cit.; Id., *Visual querying*, cit.

[8]E.g., Neo4j.

[9]W. Didimo, L. Giamminonni, G. Liotta et al., *A visual analytics system*, cit.

[10]W. Didimo, L. Grilli, G. Liotta et al., *Combining Network Visualization*, cit.

[11]*Ibidem.*

* * *

Processi decisionali efficienti e affidabili tramite analisi visuale con metodologia *human-in-the-loop*: un caso di studio sulla valutazione del rischio fiscale

Riassunto: Il *data mining* e l'intelligenza artificiale sono sempre più utilizzati nell'analisi dei dati. L'ideale sarebbe ottenere un'automazione completa, ma in molte applicazioni ciò comporta rischi significativi. In questi casi è necessario il coinvolgimento diretto di analisti umani per raffinare l'analisi o per prendere le decisioni finali. Un problema rilevante è quindi come garantire un processo decisionale efficiente e affidabile nel quale gli esseri umani sono parte integrante del processo di analisi. Proponiamo una metodologia *human-in-the-loop* che sfrutta il *data mining*, il *machine learning*, e la visualizzazione per migliorare il processo di analisi. Un elemento chiave è l'uso di una *dashboard* visuale intuitiva, di supporto all'individuazione di relazioni e pattern di dati nascosti. Come caso di studio, descriviamo un'applicazione di questa metodologia per l'analisi del rischio fiscale nell'ambito delle attività dell'Agenzia delle Entrate.

Keywords: *Human-in-the-loop* – *Data mining* – Analisi visuale – Visualizzazione di reti – Valutazione del rischio fiscale