Deep Science Publishing
https://doi.org/10.70593/978-93-49910-38-6_1

Chapter 1: The convergence of artificial intelligence, machine learning, and deep learning in tax and consulting

1.1 Introduction

The value and risks of the convergence of AI, ML, and deep learning in tax and consulting are examined. Unless barbs are in check, they can be weaponized to attack unsuspecting business clients. The monstrous Ed Trunk, a figure of rabid fear, anxiety, and paranoia about light being shone on their siloed, treasure-holding hinterland, is the first to expose the monster’s capabilities while also seeming to yield employable fact-checkers and great consultants. It should be noted that the level of deep AI will play a large part in determining who’s positively engaged and who isn’t.

The historic context of AI, the first boom and bust, is worthy of a re-examination. Inheritance can make it difficult to climb high and go low. It is quickly becoming irrelevant, but memories of it may last. Although deep learning is, in its current form, merely a copy-cat from long ago, the writer's doubt that anything like the Current Transformation will occur again.

Also examined are the legal and regulatory implications of the convergence of AI, ML, and deep learning in tax and consulting; groups of industry firms that can address many of the tax, legal, and regulatory implications of the convergence in a way that corresponds to their credibility; and recommendations of steps that tax and consulting firms should take to address the most urgent risks associated with the convergence (Gounagias et al., 2018; Savić et al., 2021;Aladebumoye, 2025). While no explicit mention is made of them, it is assumed that threats of upheaval will lead to consolidation. Therefore, moves to engender trust, as viewed through a lens, naturally lead to joint trust-related offerings, which may force some trust-oriented groupings into consideration through demands for collaboration or regulation.

Its basic premise is that the speed of AI coming into tax and consulting is greater than the industry capacity to handle or absorb it, and that there are specific actions that require greater credibility than any one firm can manage. It points to groups of firms that need or would benefit from catalyzing a working relationship with grand strategic thinking capabilities that contemplate the current amalgamation of AI, ML, and deep learning.

Fig 1.1: AI/ML development and consulting services

1.1.1. Background and Significance

Artificial intelligence (AI) and machine learning (ML) are the topics of fundamental scientific and applied interest in our fast-paced digital economy shaped by global competition. Businesses are obliged to speed up their digital transformation processes to remain competitive. Data-driven decision-making based on intelligent analytics is a key ingredient in this process. The main components in this regard are AI-enabled information management and business analytics. Business analytics converts raw data into valuable business insights. The area of business analytics is hotly debated and subject to a variety of terms. Before explaining the idea of business analytics, it is necessary to clearly understand what is meant by the terms data, data analytics, and big data.

Intelligent data-driven business analytics has been a topic of considerable scientific and applied interest during the last decade. In its broadest sense, business analytics is covered by the area of business informatics. Business informatics is an interdisciplinary field drawing and combining expertise mainly from machine learning, statistics, information systems, operations research, and management science. Practitioners have, therefore, trouble set frameworks showcasing ideas and implementations that are actually successful in practice. Ideas, applicability, risks, as well as ethical aspects of intelligent data-driven business analytics as best practices, are collected and substantiated by case studies of application areas spanning from partly to fully automated business processes.

Machine learning-enabled business applications form a great hype and promise as real success factors for enhancing competitiveness. But where are the frontiers of most success? What is machine learning anyway? What are the main ingredients of an intelligent data-driven decision-making process? What are the main hurdles and risks when implementing such applications? In response to the above questions, a machine learning-based price setting model is presented as a best practice for value creation. Different hybrid models of regression trees and deep neural networks are carefully compared. The main components of contemporary ML-enabled business analytics and information management targeting managers, decision-makers, and experts in a condensed explanatory manner are summarized.

1.2. Understanding AI, ML, and Deep Learning

The terms Artificial Intelligence (AI), Machine Learning (ML), and deep learning are often used interchangeably. However, there are subtle differences among these terms. The term AI (or "intelligent systems") refers broadly to the capability of a machine to imitate intelligent human behavior. It can be identified as the simulation of human intelligence for computer systems to perceive, reason, and act (in ways that are deemed "intelligent"). Currently, most of the intelligent systems offering AI capabilities (or claimed AI) rely on ML techniques or algorithms (Zheng et al., 2018; Zhang et al., 2022). Under the umbrella term AI with general connotation, the specific term ML is used to denote the aforementioned engineering discipline with an emphasis on the application of learning models trained by data. ML describes the capacity of systems to automatically learn from problem-specific training data to build analytical models and automate the process of analytical model building to solve the associated task of interest. A learning task is considered to be a ML task if it satisfies an objective measurement criterion and the systems rely on training data adaptable to these systems prior to deployment. Models of the learning task and generalization function ability are considered decision-making understanding/modeling according to the aforementioned characterizations.

This historic divergence in expectations poses an existential question for AI and the consulting industry itself. Governments' closing remarks on AI v. safety concerns ask, "What is our next step?" AI cries out "Why did it have to be so loud?" AI immigration reform, jobless deifying, governance preemption, and quadratic core coding policy are all grave issues in their own right, all needing explicit direction from government sources. And there are others who have satellites searching in vain for the post-human world, or the government searching for airborne clogs, bombers evading interceptors. AI, on the other hand, is itself a semantically blind god, a murder ballad no less, dividing and subjugating the peasant world while singing of a better life.

1.2.1. Definitions and Key Concepts

Machine Learning (ML) describes the capacity of systems to accept data as input and recognize potentially useful relationships and patterns to produce an analytical model as output. In particular, machine learning describes processes through which systems search for analytical models such that the best model is selected from the search space and the automatically learned model is produced. Deep Learning (DL) is a concept in ML that builds and trains networks which consist of many layers of interconnected processing units to learn representations directly from raw data. DL can automatically learn hierarchical representations in massive data from low to high abstraction levels and capture complex relationships with a high capacity of representation. Due to continuous improvements in DL techniques, computational power, and the availability of large training sets, DL is seen as a break-through technology and a true engine of data-driven intelligence. With new programming frameworks, such as Tensorflow or Pytorch, frameworks and services to run these models, as well as access to GPUs, starting at a few dollars a month, a rapidly increasing population of data scientists and developers build and deploy ML applications. Unsupervised and semi-supervised techniques reduce the required labeling effort but it still remains a barrier in some domains, most notably in the medical domain where expert labeling is expensive and time-consuming. Nevertheless, deep learning frameworks provide exceedingly easy access to pre-trained networks which can analyze a wide variety of inputs which is highly appealing to companies.

1.2.2. Historical Context and Evolution

Firms have long wrestled with actors such as artificial intelligence and machine learning. They have vague hopes and fears surrounding how these technologies will complement or replace the dreaded tax-preparer. These worries are magnified by the insights of a few fantastical and dystopian AI visionaries. There is talk of AI "doing" tax returns or giving voice to the deductibility of a child seat. AI will look at a hundred million IRS rules and regulations and a billion pages of financial transactions before announcing the findings. This scenario casts a dark shadow over millions of accounting jobs. This does not seem right.

Less cataclysmic are the relatively unglamorous, existing applications of machine learning in tax and by financial services that quietly do not help the public good. Surveys regularly reveal that sociopathic tax tech firms in the United States, UK, and Australia fruitlessly spent millions on tech to produce faster 2000-item tax returns. The results are predictable—taxes end up looking like a confusing scene where a virtuoso drummer pounds out a confusing 60 items per measure. One participant stated that lawyers nationwide would soon have to readjust their format from “Two weeks, half a million for the privilege” to “Go to junior college, one day for $99.95” as AI replicates encroaching tax law on drones forwarding cousin Andy’s barbeque dinner dates. Even negative outcomes such as compliance issues warrant interminable and futile discussions about early AI work.

1.3. The Role of AI in Taxation

The latest multi-disciplinary research into artificial intelligence (AI) brings together scholars from machine learning, computer science, mathematics, economics, and philosophy to explore the potential for transformative advances in the future and the implications for humanity. There is a nascent, but urgent, call for self-regulation among AI/ML (machine learning) technology companies and for some form of regulation of AI by governments or intergovernmental organizations. To date, regulation, oversight, and governance of AI have focused mainly on ethical issues from a legal perspective. AI is defined, limited, and controlled by legislation developed by existing political institutions concerned with the implications of technology. The goal is to slow technological advancement and to ensure equitable use of developed systems. While necessary, these plans may miss a key nuance: models themselves are not inherently better or worse for society. Technology can be molded to one’s devices, often in surprising ways. As technology becomes enmeshed in society, it must adapt to accommodate new realities and capabilities. Legal developments must acknowledge this nuance and ask the opposite question: how can AI advance political goals—in net beneficial ways? It is broadly accepted that AI has the potential to have a huge and positive socio-economic impact. As historical precedent, the introduction of railways led to economic growth, people's movement towards urban areas, and new locomotive-dominated cultural phenomenons. These changes pushed innovations in regulation, transportation systems, other historic events, proxies, etc., that brought modern-day cities, economies, nations, cultures, and societies. In accounting and consulting, continual shifts in tech, regulations, and society have created new needs for firms around transparency, safety, trust, and understanding of actions and predictions. Transparency is often captured by data and algorithmic accountability and 'right to explanation'. It often boils down to either rule-based or post-hoc explanations. Interpretability can refer to the ease of use, understanding, and trustworthiness of ML predictions and actions. IML also refers to the renewed focus of ML research on 'wide' models trained on structured and semi-structured data outputs, their use of interpretable probabilities, feature distributions, and properties, and interaction with increasingly powerful reasoning. Regulation of ML applications has often been a reactive cat and mouse game, with tech advancing faster than legislation, and with the understanding of tech often lagging behind its actual implementation and consequences, especially for corporate technologies that researchers have little access to. Equitability is today often defined by conformity to societal norms while uncovering unintended goals that can empower the status quo.

Fig 1.2: Artificial Intelligence in Taxation

1.3.1. Automating Tax Compliance

The IRS has long required businesses to complete and file tax forms based on information returns prepared and submitted by third parties. Various information returns are required depending on whether an item is real or personal property, whether it is paid to one or more individuals, whether the payment is for rent, royalties, or dividends, and so on. For example, for landlord–tenant relationships involving rental payments, Form 1099-Misc and Schedule E must be filed with an IRS issuance in the following tax season. Similar reporting requirements at the state level will vary widely and are not currently available in the form of tax models. The same goes for K-1s and state S-corp tax returns. Tax compliance becomes a multi-entity process, with five or more tax forms exchanged among clients, tax returns, and consulting firms. The need for such tax models within and outside the United States creates opportunities for the tax and consulting firms.

Whole fields of tax and consulting, client entities, and tax documents are needed in order to automate tax compliance. Entity types include but are not limited to individuals, corporations, partnerships, estates and trusts, and non-profit organizations. Entity subclasses include narrow casts of company sectors such as real estate, construction, and retail firms. Tax documents are those forms, schedules, and declarations filed with the IRS and with state revenue agencies, classified into five groups of various issuers: individuals, partnerships, corporations, estates and trusts, and exempt organizations. Each of the five groups may include at least a dozen tax forms. Unfortunately, there are currently no tax models for the tax-preparation tasks and before-tax scenarios just mentioned. The default evidence framework is the predecessor of the private ChatGPT. The first version of this product can extract and summarize financial documents at the level of items, totals, and types.

Completing tax forms is based on an understanding of the laws and principles upon which the forms are built. Therefore, classical logic will continue to underpin generative AI. Nonetheless, using black-box neural networks to represent a company's global knowledge base is unprecedented. The availability of open-source large language models will help conduct experiments on crafting the knowledge base into vectors stored in a vector database . Possible implementation steps include querying laws, rules, and regulations on-demand via APIs, rapidly transforming themselves with machine-generated feedback, and cross-referencing names and numbers of the forms relative to the knowledge base. New LLM zero-shot and few-shot capabilities to understand tax implications on real-world computation across various tax rounds will allow for natural language tax-generating software. Namely, tight integration of tax systems and open-source LLMs will boost the productivity of existing human resources without redundancy, making themselves even more cost-effective to improve completions and turning the software into multilingual and multiregional tax expertise out-of-the-box, as well as straight-through processing of tax sources irrespective of formats.

1.3.2. Enhancing Tax Planning Strategies

One of the primary areas that the merger of AI, ML, and deep learning will impact tax and consulting is enhancing tax planning strategies. The legal acceptance of these models allows for wider adoption at places like the Big 4 accounting firms, which may be forced to develop their own models or incur large costs to use. Early returns from publicly available models have been promising in generic use cases covering things like legislative explanation writing style, suggest filtering, and accounting principle applicability. In tax consulting and tax planning strategy modeling generative question answering tasks, there are challenging aspects such as solving question relevance found in different sub-regulations and legal principles as well logic representation challenges. Nevertheless, these tasks have enormous social value.

The questions and user requirements presented to the model allow for some differences to be expected. The fitness expectancy of the model could be decided by mentioning the judgment perspectives including discussion basis, model type, timeline, evaluation criteria, and relevance discovery method in addition to the questions about tax consulting. To answer questions regarding tax planning strategies, the relevant knowledge mainly contains the commercial tax planning strategy and the related law. It could be retrieved using the keyword-based method as there are strict terms in the relating acts. However, due to the discrimination of legal text understanding and prediction from regular language modeling tasks, the default relevant knowledge retrieval method from generic text LLMs plus vector databases might have relatively limited performance. Hence, this study focused more on text choosing than retrieval.

Unlike the broader practice to encapsulate more explicit tax domain structures, the focus of this study is utilizing the precise commercial law and targeted question construction to challenge LLMs for generating guidance about tax planning strategies. Retrieving knowledge of tax law is tackled using a KFS based text filtering method that establishes a ballpark to be considered. Then Q/A pairs generation is tackled through best prompting and decoding methods searching, similar to interfaces optimizing kind, which are highly sensitive to questions and require adjustment of models. The retrieval models' possible success and effectiveness are examined under different conditions, such as the combining of tax law text configuration and filtering structures. In addition, guidance generation for tax planning strategies is executed through thoroughly chosen control constructs and template engineers. To benefit more public models, all the parameters and hyperparameters are tried to be shared, like the prompt and text selection criteria. Overall, during this positive cooperation of modeling text representation and procedural understanding, the creativity of legal capability may grow as shown in the meticulous cases.

1.4. Machine Learning Applications in Consulting

Machine learning (ML) is a subfield of artificial intelligence (AI) where machines learn automatically or from experience. This learning can be either supervised or unsupervised. Machine learning requires an extreme amount of data to train and study the model. Consulting is a growing industry worldwide. When consulting firms face big data, it becomes hard for them to offer an accurate solution. It also requires a huge manpower force and time to evaluate large data. Machine learning can be used to overcome this hurdle and can offer faster and dynamic results using automatic learning models. Machine learning has numerous applications in the consulting industry.

In the consulting industry, data is collected from different sources and then cleaned as per requirement. After that, various models can be created. Very often a firm gets a huge amount of data for consulting. It is impossible for consultants to study that large amount of data within a given time. The tagging of data will also take a huge amount of time. If the data are untagged, then there is no possible way to create a model. Here, unsupervised model creation can work. By creating the model, the original data can be studied more efficiently, and the pattern of data can also be plotted. The patterns help in deciding according to the requirements of the problem.

For consulting purposes, the output of study reports that are expressed in natural language is required very often. Consulting firms use analysts for offering reports who study original data nodes or interim results. A better alternative can be created by using models. By using training templates of reports written by analysts and associating them with nodes in the models, industry standard report formats can be automatically created. The model also examines every connection of the original data nodes and interim results nodes with these report templates and selects the most applicable ones.

1.4.1. Predictive Analytics for Client Insights

During the tax compliance process, 100 percent of a client's data is typically acquired and reviewed to assess the risks involved for that client. Within this workflow, the attention of the tax team largely follows the principle of Pareto. In a tax return, 80 to 90 percent of the risk typically is concentrated in 20 to 10 percent of all transactions. Accordingly, to effectively manage client risk, it is crucial to accurately identify the client's riskiest transactions beforehand. Predictive analytics makes it possible to assign a risk score to each transaction with the help of ML as a basis for checking the tax returns.

For businesses, decisions in areas such as accounting and finance can greatly affect their ability to compete, enter new markets, reshape product portfolios, and prioritize research and development. The ever-changing views of tax administrations regarding acceptable and unacceptable transfer pricing practices create an urgent need for businesses to proactively assess their compliance risks and exposure instead of reacting to inquiries and examinations. The rapidly growing transaction-level data sets and the increasing sophistication of verifiable yet unstructured data represent an opportunity for businesses to adopt fully automated compliance processes and advanced analytics to develop new compliance insights and processes and enhance compliance in a cost-efficient way.

Today's tools for developing analytically quantifiable compliance insights are largely focused around query/buttons/computational visualizations with descriptive/or ad hoc analytics capabilities. They are disconnected from a business's tax compliance workflows, ease of use, and versatility, resulting in practical barriers to entry, adoption, and application of advanced analytics. Further, the complexities involved in assessing compliance risks call for comprehensive and holistic solutions. Complex compliance-related questions often cut across multiple datasets and competencies. Full-scale automation of compliance processes often requires deep integration of analytics with transactional bookkeeping ledgers and middleware.

1.4.2. Risk Assessment and Management

Because sophisticated adversaries—both foreign and domestic, state-sponsored and independent—pose a threat to the stability of the nation's financial systems, significant resources are dedicated to the hardening of these systems. Regulatory compliance is a general requirement across industries in the United States, ensuring cybersecurity for both investors and consumers. The general definition of risk, as implemented in financial services, refers to a subset of threats that come from motivated adversaries who are attempting to exploit system weaknesses to cause direct or indirect harm—be it economic, reputation-based, physically destructive, etc. As these economic threats evolve, risk management and compliance systems must remain vigilant, updating their training data, alert methodologies, risks prioritization, mitigation techniques, and detection thresholds. Increasingly complex threats can rapidly advance past the capabilities of trained systems, actively seeking to obfuscate their methodologies and exploit platform weaknesses before traditional governance frameworks are able to adapt. Financial institutions across the globe are relying on Deep Learning (DL) technology to implement timely and sophisticated risk assessments and active mitigation measures to combat evolving threats.

While traditional financial risk diligence can provide a snapshot of a single firm's relative position with respect to reported balance sheet metrics, a deep learning, domain-specific knowledge graph of financial markets can be constructed using unstructured sources of disclosures, news articles, social media, etc. A key area of sentiment and event analysis is the regression of quantitative market reactions using semantic features. Information from entire articles, messages, and transcripts can be monitored in real time and sorted for the most salient events. This can help detect misleading disclosures, which might form the basis for criminal insider trading, as well as shape more granular assets-level sentiment analysis. Anomalous changes in predicted asset value movements provide a basis for proactive escalation of alerts to analysts, who can then verify potential Governance, Risk, and Compliance (GRC) issues and trigger event-specific monitoring. Anomaly monitoring is especially useful given the typical temporal nature of risks (short-lived events filtered by rolling windows), as it requires only examining recent data. Traditional classification models can be trained using historic data to identify these low-level governance risks.

Despite significant developments in various areas of AI and diligence, there remains a notable absence of tools that allow for effective, standalone risk assessment of agent-based systems. Current governance frameworks are inherently qualitative: organizations often come to learn of a deployed model's effects post hoc. These updated recommendations ensure responsible design by teams empowered to make modifiers on architect-trajectories less than ideal for currently misbehaving legacy models. Use cases include explicit ubiquitous hiring risk-disclosures that would be attachable to any job-ads enhancing hiring requests with no restrictions on disclosure/design teams. Quantitative diligence is limited to actometrically calibrating edge-case robustness. With pilots proposing to impose minimum safety deadlines without determining what they apply to and how to enforce them, large language models and image generation systems pose unprecedented open cases.

1.5. Deep Learning Techniques in Financial Analysis

As contextualized in previous chapters, Machine learning (ML), a subset of Artificial Intelligence (AI), refers to statistical techniques that learn the representation, patterns, and relationships in the data. Artificial intelligence (AI) automata or mechanisms that frame rational behavior are based on a hard-and-software system attempting to exact particular types of rational behavior. "Deep learning" is a noteworthy field of computer science that studies architectures, algorithms, models, and frameworks with multiple levels of processing. The convergence between AI, deep learning, and ML is evident with the emerging new players that are activating Automation, Robotics, and AI by incorporating deep learning, AI, and ML systems. The convergence can improve productivity, efficiency, accuracy, and cost. The discussion is based on the most recent applications of deep learning in accounting, auditing, assurance, tax, and consulting.

Financial analysis assists investors, financial analysts and managers in generating objective and insightful investment advice or financial judgment. However, oil prices are influenced not only by historical charts, trends, and cyclical factors, but also by time-stamped, discursive, complex, and unstructured textual economic indicators. Hence, the challenges and concerns for financial analysis as a sequence of the digital transformation wave are enormous. Financial disclosures with report characteristics, such as MD&A, numbers, length, rate of growth, and the sentiment embedded in text can be collected and subjected to analysis. These enhancements empower financial analysts and investors to predict stock price movements. In an unprecedented volume, financial disclosures are created to present, convey, and demonstrate. A pressing yet astonishing issue is whether the vast and complicated financial disclosures can suggest the variability of capital market reactions.In the last two decades, methods of machine learning have been extensively researched and papered on in order to tackle those increasing data granulation and frequency. However, most machine learning methods suggested are still based on the classical linear structure assumption, making it often unsuitable outputs or forecasts for nonlinear financial data. Furthermore, the sheer complexity of the financial markets with a myriad of inter-relation makes analyzing it in the big picture through model fitting hard to deal with.

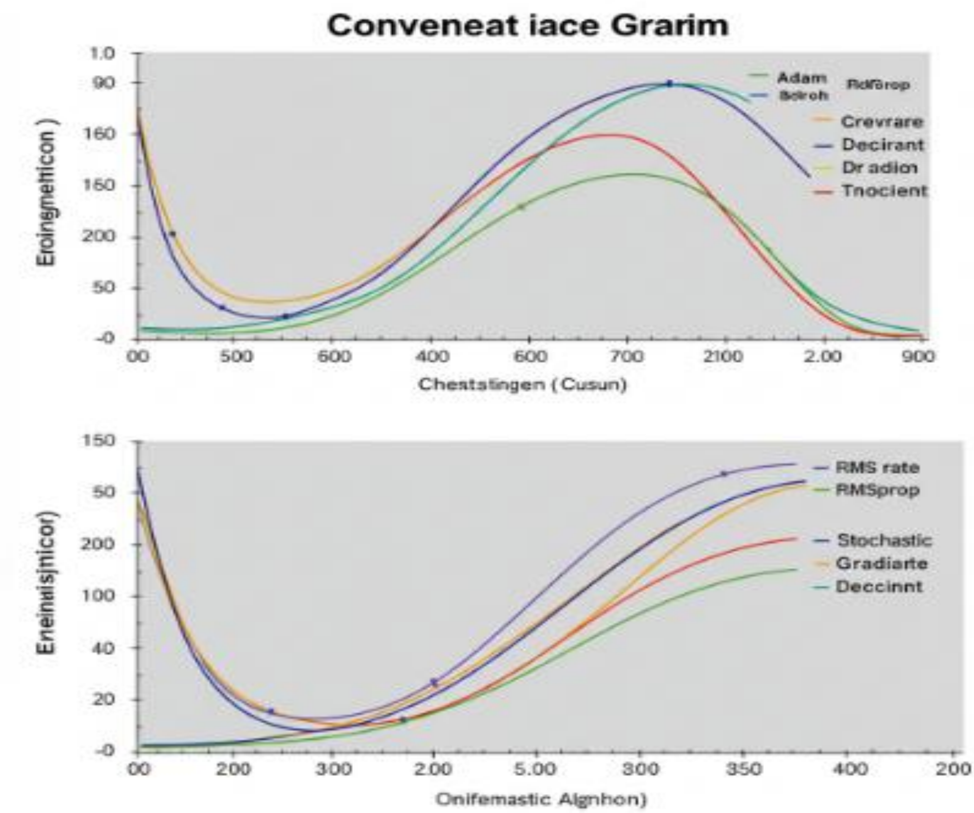

Fig : Convergence curve of each algorithm

1.5.1. Neural Networks for Financial Forecasting

Artificial intelligence is becoming a crucial tool for immediate analysis of financial data as economies shift from cash to digital methods . However, deep learning, which is a form of machine learning, has mostly only been explored in a few academic finance topics areas, and fewer successful implementations in the finance industry are seen. Popular research in finance and deep learning areas include classifying stock prices movements, predicting stock returns, and using social media or news data to predict stock returns.

It is known that financial markets are complex systems that are affected and driven by a multitude of independent variables, many of which themselves change over time. Therefore, financial data has characteristics such as chaos, changing, seasonality, volatility, noise, and non-linearity. In recent decades however, advances in computations have led to the considerably increase of analysis methods for the financial data from a statistical point of views. At the same time, the big data problem when analyzing financial data has become an issue as the stocks, indices, and commodities of interest grow.

1.5.2. Image Recognition in Document Processing

The application of image analytics in the review of legal documents aims to automate the time-consuming and expensive process of screening through image documents that support e-discovery. A technology-assisted image review pipeline is introduced that uses state-of-the-art deep learning techniques for document-level image classification, image clustering, and object-level object detection applications. By leveraging transfer learning techniques on established pre-trained models, good accuracy is achieved while reducing amounts of training data and time needed to pick up computer vision techniques. Applications of deep learning in computer vision to enhance legal technology-assisted review (TAR) are presented. Prior TAR applications in the legal industry only deal with text documents. In today's litigation context, many documents are images, and they often play a critical role in summarizing important information about cases. For example, in e-discovery, an organization is typically required to produce old documents that are related to cases and image documents such as scans of old paper documents and screenshots are common. However, due to technological challenges and the unstructured nature of those documents, manual review is the norm, making the review more time consuming and expensive. Several business cases of image document review projects in the legal industry are studied and their pain points are analyzed. The legal document review process requires transforming image documents into a structured format for further analysis. Although some generic image processing techniques can be applied before legal review, most techniques have limits with regards to different sources and formats for image documents. There is thus a strong demand to develop customized applications on image data. Beyond the applications focusing on documents, items, scenes, and other contents depicted in the image, extensive image tracking for scene and character animation visualization in the media industry is invoked. Further, the accomplishment of detecting business card boxes indicates the feasibility of applying object detection techniques for other images data.

1.6. Data Privacy and Ethical Considerations

There are a number of privacy risks and ethical considerations around the usage of personal and sensitive data in tax and consulting services. The intersections of AI, ML and Deep Learning deliver important insights for businesses and professionals, including Tax compliance detection, Tax risk forecasting and value creation, Consulting engagement detection, and result explanation. However, the results also carry serious privacy and ethical risks. The main sources of such risks are data privacy and ethical considerations. They must be reviewed when professionals are provided with insights based on AI/ML/Deep Learning models.

There is a strong need, academic scrutiny and industrial demand on understanding the intersection of data privacy, ethics and AI/ML/Deep Learning models. The directives, regulations and seals around data privacy protection are normal around the world. Examples are the General Data Protection Regulation of the EU, the Children’s Online Privacy Protection Act of the US, the Personal Information Protection and Electronic Documents Act of Canada, China’s Personal Information Protection Act, similar efforts in the UK, Brazil and Singapore, as well as Privacy Shield framework and Cross-Border Privacy Rules adopted in 5 continents. Therefore the incoming business models based on the proposed methods are on the verge of being conducted with potential huge long-term competitors’ advantage.

The provision of tax compliance inspection results in such new methods analyzing a huge scope of sensitive data exchange, leading to huge concerns on data privacy violations, ethical discussions and possible misconducts. On the other hand, even though taxation and consulting services are nation commanding business, they remain a lack of academic scrutiny, meanwhile being less regulated under the perspective of data privacy protection. Even new media and proposed methods attracted some attention in general supervision frameworks, the domain consultation and taxation data modeled in tax archivers and numerous business consultants poses a critical component of a nation.

1.6.1. Regulatory Compliance in AI Usage

The long list of AI governance challenges motivates the need for new solution approaches. Central to an AI system's design and deployment without AI model governance is a new approach to continuous regulatory compliance in AI usage. The term "governance" is used to encompass this broader focus, while AI model governance or "model governance" is used to specifically refer to governance considerations that solely focus on the AI models. An AI regulation and compliance framework for the financial services ecosystem is proposed, built upon the classic governance systems architecture. The proposed framework lies at the boundary of the AI and governance systems. Continuous regulatory monitoring and reporting during deployment shows the high-level architecture of the proposed AI system framework and regulatory model. The governance systems in the proposed framework aim to incorporate run-time monitoring, regulatory control, and mitigation capabilities in the production environment.

As a broader focus, AI governance can take many forms and operate on various levels, including firm-wide governance and macro-prudential governance. Meanwhile, the word "model" is used broadly to include all types of ML algorithms or AI models, structured or unstructured inputs, traditional or advanced techniques. The framing assumes that the model itself is a black-box representation, so that its inner workings are pre-processed away by means of some transformation such that the inputs, outputs and other governing dimensions, like fairness and robustness/scalability, can no longer be reconstructed from the new representation, but can be monitored instead. AI algorithm is also not limited to but includes ML or statistical models compatible with such preprocessing, as most of the AI governance challenges arise in the finance industry's use of ML, and this framing intrinsically connects with interpretability enhancing techniques.

1.6.2. Ethical Implications of Automated Decisions

The usage of artificial intelligence, machine learning, and deep learning in tax decidable decisions in a business environment raises challenging ethical issues. These are decisions made by computer programs or algorithms that may use some data, but the true mechanism or logic behind how the information was processed is not known to human users. This is often referred to as a "black box" problem. There are wide ranging critical concerns regarding systems that automate decisions that are not transparent and that creators/users cannot explain how they work or why they arrive at a particular conclusion. Such decision systems may violate the basic requirement of accountability.

The ethical concerns regarding black box decisions fall into three broad categories: (1) advanced technologies or systems that enable black box decisions; (2) decisions that depend on the technologies; and (3) the safeguards necessary to prevent abusive or unreasonable decisions. The notion of automatic decisions being made by computer programs or algorithms that are too complex, multilayered, or non intuitive for even their developers to understand is often a key concern regarding advanced technologies. Examples include technical descriptions of Artificial Intelligence (AI) applications such as machine learning (ML) and the different dimensions of this space, neural networks of different types or variants, etc. Beyond the term “algorithm”, it is even more difficult to conceptualize what a “Deep Neural Net” (DNN) is .

Having descriptions or a high-level understanding of how systems are created or how they function, the worry is that they are run by a combination of rules, weights, parameters, and training data that no human could expect to interpret or comprehend.

Independent of the internal structure or mathematics utilized, there is a wide range of decisions — of variable types and computational complexity — that might render human explanation a challenging task. Others yield a representation where there are overwhelming character limits. Furthermore, no system is perfect, and there are events in which a false decision is made, which resulted in harm.

1.7. Conclusion

AI has started to impact many aspects of life these days. Artificial Intelligence (AI) is the field of study that creates intelligent agents. Those agents can be physical or virtual, and can be made of hardware and software or of one of both. In general, AI consists of algorithms which calculate the states of its concern and give commands to affect its environment. Using its sensors, an agent collects data from the environment. Then, it can estimate parameters of interest using algorithms to give meaning to the collected data. Finally, it can interact with the environment by executing behaviors and actions through its actuators. It’s basically a loop of perception, cognition, and action to achieve system specific objectives. AI is a wide field of study. There are many aspects of it. Yet, this text focuses on the recent advancement of a particular subfield of AI which is known as machine learning (ML). Machine Learning (ML) is the subfield of AI that studies the creation of software capable of autoregulating its behavior based on the input data. Generally speaking, these algorithms can generate models to process input data and obtain predictions out of it. In doing this, the algorithms have two phases. The first phase is called learning. During the learning phase (the training phase), the software takes a training set comprising input data and the expected output. With this data, the software creates a machine learning model containing a set of internal parameters that will process the information later inside the second stage. The results of this phase are not yet predictions. In the second phase, referred to as the inference phase (the testing phase), the software takes a new data set to be tested. This new data contains only input data with no labels. Using the model created in the previous step, these inputs are going to be processed in order to obtain the predicted outputs.

1.7.1. Future Trends

Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) continue to progress rapidly. As this evolution occurs, the key challenge in nearly every industry is how to use the opportunity to build smarter organizations. Each of these technologies is mature enough now that organizations are starting live cases. However, some organizations are progressing more than others. This framework can assess which stage an organization is in with respect to digitization opportunities.

The first stage is defined as digitization. This is the easy part: organizations use digital sensors such as RFID or imaging to record digital fingerprints of products, transactions, documents, or events. Data ownership changes forever when data is recorded as zeros and ones, which can be replicated infinitely, stored indefinitely, and analyzed cheaply. It is difficult to argue these changes don't significantly alter the competitive landscape of these industries.

At the second stage, organizations analyze data to determine what happened. They use manageable SQL databases, simple tabulations, or 2-dimensional data visualizations to demonstrate problems. Consider a coke bottler operating in a thousand-restaurant chain with a few thousand suppliers. Performance is generally viewed by checking how many units of coke each restaurant sold analyzed as a 2-dimensional line chart. Instead, organizations should analyze their market share across restaurants, products, dispensing frequency, competition intensity, economic conditions, holidays, nearby banks, growth, and testimonials using larger multi-dimensional product family databases.

References

Aladebumoye, T. (2025). *Integrating AI-Driven Tax Technology into Business Strategy*. International Journal of Research and Review, 12(1), 224–231. IJRR Journal

Gounagias, N. D., Hristu-Varsakelis, D., & Assael, Y. M. (2018). *Using deep Q-learning to understand the tax evasion behavior of risk-averse firms*. arXiv preprint arXiv:1801.09466. arXiv

Savić, M., Atanasijević, J., Jakovetić, D., & Krejić, N. (2021). *Tax Evasion Risk Management Using a Hybrid Unsupervised Outlier Detection Method*. arXiv preprint arXiv:2103.01033. arXiv

Zhang, Z., Al Hamadi, H., Damiani, E., Yeun, C. Y., & Taher, F. (2022). *Explainable Artificial Intelligence Applications in Cyber Security: State-of-the-Art in Research*. arXiv preprint arXiv:2208.14937. arXiv

Zheng, X., Zhu, M., Li, Q., Chen, C., & Tan, Y. (2018). *FinBrain: When Finance Meets AI 2.0*. arXiv preprint arXiv:1808.08497.