Deep Science Publishing
https://doi.org/10.70593/978-81-988918-1-5_4

DeepScience
Open Access Books

Chapter 4: How generative artificial intelligence is revolutionizing business intelligence, creative processes, and law

4.1. Introduction

Generative Language Models (GLM) work by predicting the next word in a sequence of words based on previous tokens. In order to perform this word prediction task, they are trained on large amounts of text data. Training involves two key components: supervised pre-training and unsupervised fine-tuning. During the pre-training phase, the models are trained using maximum likelihood estimation to minimize the difference between true and predicted tokens. For fine-tuning, human annotators provide guidance by evaluating model responses to prompts, while reinforcement learning is employed to improve model performance. Supervised pre-training enables the models to leverage vast amounts of unlabeled data, while fine-tuning ensures that user intentions are taken into account. This two-part training procedure is similar to that of animal learning, where the first phase features unsupervised learning to form a general model and the second phase involves reinforcement learning to make the model comply with specific user instructions (Bommasani et al., 2021; Liu et al., 2022; Goel et al., 2023).

To enable GLM capabilities, two major considerations are required. First, training requires massive data and compute. Existing models have 10 to 1,000 times more parameters than earlier models, and training of a single model may be accomplished with 1,000 to 1,000,000 GPU days. The recently proposed Mixture of Expert architecture uses a smaller number of activated parameters during selection than the total number of parameters during inference, leading to efficiency. Second, scaling laws for model size, data size, context length, and compute must be respected. A trained model performs better when the training set is larger; more gradient updates are performed; more parameters are used; more tokens are used from the pre-training dataset during each

update at each scale; for larger context lengths; and for denser pre-training on a larger bucketed pre-training dataset (Liu & Jatowt, 2023; McCormack, 2023).

4.2. The Evolution of Business Intelligence

The business intelligence field has been experiencing an amazing and rapid growth. Recently, we have witnessed an increasing interest for technologies that can automatically provide analytic insights to non-technical users and business decision makers. We may define these technologies as business intelligence natural language interfaces. The reason for this development is twofold. On the one side, tech-savvy

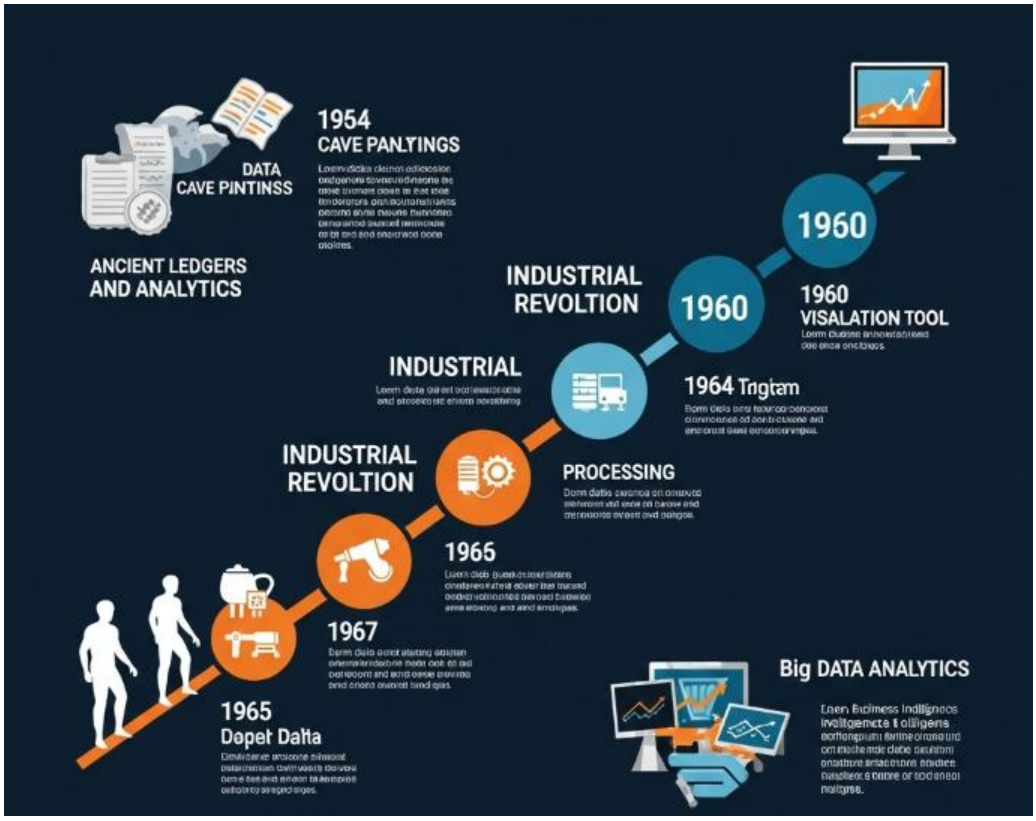

Fig 1 : The Evolution of Business Intelligence

people, like analysts, computer scientists or information system developers, understand that data analysis is a crucial part of the decision making process in any modern business and that by making this process easier and faster, they help their company, or their customers, achieve their goals. On the other side, the people that rely on the information provided by these systems are not technical users, they are very often business representatives that want to obtain analytic insights based on their knowledge of the

business problems or opportunities, but that do not have the technical knowledge required to interact directly with the underlying systems.

The need for tools and services that allow these business experts to obtain insights directly from the data without relying on technical analysts has been present for several years. The actual goal is to provide these people with systems that allow business decision makers to access data, retrieve information and obtain report-like outputs for specific questions, without relying on the causes and additional specification that are required, traditionally, from BI front-end tools. In order to achieve this goal using BI NLI Systems, experts need to express their questions using concepts that are more or less familiar to them, and these systems need to properly translate their expressions into the analytic tasks that lead to the generation of the desired results. Business Intelligence Natural Language Interfaces have been designed for data access and data analytics by non-technical users. BI NLI Systems are based on the possibility of creating flexible data models and are capable of supporting different NLU models and designs.

4.2.1. Historical Context

In the early 90s, business intelligence meant searching and gathering information. In the early 2000s, the term began to mean the analysis of gathered information. Nowadays, business intelligence primarily means the distribution of reports that present key performance indicators for better enterprise management or self-service, in which undiscriminating users can seek their own answers. The BI segment has been one of the most responsible for the growth of the software industry in the last decade. As other sectors, it has grown and matured. During the years, it has opened up new opportunities and given rise to new players. The founding of this Association in 1989, and the successive fundings of other related associations around the world, more than a trend, are concrete facts that demonstrate the existence of a sector and a vendor community that is more than prepared to cope with any near future demand for data warehouses, data marts, OLAP servers, Alerts and Scoring engines and analytical CRMs. There are companies assessing and implementing data warehousing and OLAP servers inbound solutions.

Business Intelligence has always been important in serious corporative environments. Yet, the establishment of the European and American stock markets, which imposed the need for firms to report their figures every trimester, showed that it was more than a fad for big corporations. Even for companies out of the stock markets, generally a lot smaller, objective budgeting and management processes were key-motivated by banks and other financial institutions. Internally, these requirements were driven from the top-down by owners and founding partners, who carried out control management thanks to software that allowed them to follow the historical and prediction figures. However, the

advent of Windows and the democratization of microcomputers, and especially the advent of client-server environments for enterprise management software, were decisive factors for the appearance of Business Intelligence in medium-sized companies.

4.2.2. Current Trends

Recent advancements in technology – particularly from Cloud computing, the Internet of Things (IoT), and AI – have impacted the development of the Business Intelligence domain. In terms of data management and operational BI, notable trends changing how data is collected and provided are the use of the Edge for data collection, the use of data lakes for mass-data storage and management, the implementation of streaming and complex event processing, APIs and data virtualization for delivering data to BI tools, the use of Natural User Interfaces, and augmented analytics. These recent advances made in analytics and analytic systems consist essentially of the implementation of predictive and prescriptive functionality in technologies such as Data Warehousing, Data Mining, In-memory computing, Data Lake analytics, and Business Analytics platforms.

Big Data is usually used to describe a phenomenon where humongous, heterogeneous, high-velocity data is collected and often analyzed for Business Intelligence providing modern organizations with competitive advantages. While companies have been using Business Intelligence techniques ever since the advent of computing, the recent data explosion has made it quite favorable to push organizations to use Big Data and analytics to implement more established processes but for new domains like Data Governance and the adoption of analytic systems such as Business Analytics platforms, Data Warehousing, Data Lakes, etc. Also, modern organizations are using Big Data analytics and AI more for financial fraud detection and geography and demographic related domains as well as risk and revenue management and market and customer segmentation.

4.3. Impact of Generative AI on Business Intelligence

Business intelligence enables the analysis of data. It includes data analysis tools such as machine learning for data analytics, predictive analytics, machine learning generated text reports from data, data visualizations. Generative AI can improve all these areas by improving code generation, making better predictions using generative models, generating reports and visualizations, or simplifying the inputs to the machine learning models and other functions. These improvements are made possible by several factors: the systems are great at generating code for these tasks, the generative models are showing great promise for making better predictions than other architectures, and the neural models are allowing for a simplification or replacement of the pre-training of

classical machine learning approaches. Generative AI also accelerates faster incremental adaptations of a method that enforces planning by gradually making the model perform as a planner.

Fig 2 : Impact of Generative AI on Business Intelligence

Generative AI makes it easier and faster for data analysts to clean data, fix problems, accelerate feature generation for predictive models, and help create default models, particularly through the better use of Natural Language Processing. A number of tools are leveraging these capabilities and showing good traction. These tools provide a significant improvement based on the ability of these systems to make use of Natural Language Processing to allow analysts to accelerate all stages of data preparation and accelerate feature generation for predictive models.

Generative AI can greatly enhance predictive analytics by speeding up the development of predictive analytics models, making them more understandable, and allowing faster updates to any prediction. Generative AI can replace traditional root-cause analysis that recommends actions to solve business problems. Generative AI can replace root-cause

analysis because a foundation model predicts the business problem variable, using past data as input. The recommendations follow from the prediction probabilities.

4.3.1. Data Analysis Enhancements

Generative AI makes data analysis more efficient. Data analysis serves as the foundation of business intelligence work, and improving this function will have a tremendous ripple effect on the entire BI process. Basically, when working with a traditional BI department, you would submit your question, and they would run a report. Then they and their data scientists would analyze it, looking at it from the lenses of every principle we've discussed so far—rules, outliers, correlations, interdependencies, clusters—then report back with their findings and recommendations, sometimes in lengthy presentations.

With Generative AI, these cycles of discovery can be done much more quickly, if not in real-time, with the added benefit of augmented version offering, guidance, and reduction of gaps in knowledge. This is a monumental shift in the "data person as sage" model of BI. Essentially, there would no longer be an intermediary between nonexpert users and the data. If you wanted to know trends for a product, you wouldn't need to ask the data team to spend a few hours running a report, and then interpreting the results. You could simply query the database yourself, using natural language to ask for data presented in the way you like best. Which in this case might be charts. In the course of the analysis, the AI might discover something unexpected. It could use terms like "surprisingly" or "unexpected" to act like a guide, steering you towards areas of interest. If it found any phenomena worth investigating—outliers, or breaking trends—it could offer them up to you.

4.3.2. Predictive Analytics

Despite the fact that predictive algorithms have played a central role in business intelligence (BI) systems for decades, providing customers with accessible data pipelines, generative artificial intelligence (AI) opens up new paths to creating predictions. Predictive analytics provide customers with predictions about what will happen in the future based on the knowledge that has been condensed into the large language models (LLMs) during training. Movie studios tell the AI prompt to summarize states what the company has done over the prior years with coefficients asked for in a recent box office prediction; the LLM makes this prediction comment based on examples it has of prior LLM-accomplished requests.

BI customers already have evidence generated by LLMs that predictions are possible from LLMs because the AI can generate predictions from prompts that ask for a

“prediction” on particular dates and times, with discussion about the “accuracy” of (a better default word on these conditions would be “validity”) is written about near the end of this section. For example, prompting the AI with instructions about every time the bulls had above pair of wins around date x and summarizing the teams, factors in charge on date x, asking for a discussion of the local factors affecting the outcome could lead to probabilities about the Bulls beating a particular team in Milwaukee – or any location for the game.

Predictive analytics, though, gets deeper than simply receiving probabilities from your model. It allows companies to ask exploratory questions about what do flies into an outcome or set of outcomes, catalog, and discuss these possible flights, and then organize them into customer-friendly packages that allow customers to ask other deep questions – how would changing x change the probability of getting outcome y (the what if)? Wouldin zi be impactful in getting y outcome?” A core goal of BI is to be able to display predicted metrics and topics in a format that business managers would naturally explore.

4.3.3. Automated Reporting

Reports have processed and presented insights discovered in data analysis. Creating reports, however, is not only time consuming; it is a rote function that is largely uninspiring, often repetitive, and does not require deep human intelligence or creativity. It may be hoped that Generative AI can help writers create better reports quickly and easily. We are already starting to see such systems come to market: tools generate analytics briefs directly in English so that readers can access customized insights without needing the expertise to wield more advanced systems. Other systems summarize analytics in natural language. Text features in such tools combine analytics capabilities with Generative AI language tools to create and enhance narratives based on uploaded data.

Generative AI has taken analysis automation to a new level with the ability to produce messages based on real-time events and analytics in language that fits a company's particular culture. Businesses can use it to enrich and quicken announcements of sales, earnings, and product-related events. Generative AI models convert statistics into sound bites while preserving sentiment, emphasis, and cadence, taking advantage of style patterns found in a company's previous announcements. Generative AI allows for advanced customization of reports. Companies can set up the Generative AI system to answer requests for specific aspects of a report. Businesses can build Generative AI engines using a combination of multiple automated reporting and data infrastructure platforms, and several companies are integrating such tools into industry-specific cloud platforms. In summary, Generative AI-based automated reporting makes the job of data reporting easier and faster, using the model's knowledge to summarize information.

4.4. Generative AI in Creative Processes

Generative AI models can create entirely new and unique pieces of content based on their training data. Works of literature, news articles, video scripts, motion pictures, video games and their assets, code, and many other text, image, audio, and video types of data are increasingly being synthesized by these systems. People are using generative AI models to draft everything from articles and reports to songs. Often, the drafts have enough polish—e.g. coherence, grammaticality, fluency, and relevancy—to warrant go-public work.

Designers and developers are only beginning to use generative AI to assist in their work; however, it has the potential to alter many creative processes. Designers have, for years, used files and libraries of other peoples' content to quickly store and access knowledge of common graphic components—from search engines to AI-enhanced engines to pattern-based libraries. Technology companies are exploring how to augment and accelerate design work in rapid iteration; for example, some have begun introducing generative capabilities within their existing products, as have other design tools.

Generative AI has seemed to come out of nowhere in creating music and visual art; that is both a cause for alarm and astonishment. Using various tools, users can generate photorealistic or artistic art in seconds. Whether or not the generated work is creative, as per Western philosophical convention, is a question as old as ancient Greece.

4.4.1. AI in Content Creation

One of the prime areas where Generative AI is making a mark is content generation, be it generally, like reports or FAQs, or domain-specific, like news creation, story writing, or academic article assistance. The new tools can produce human-like text at scale. The initial text offered also can be without too many reservations of style or substance, and the results can be engaging and valuable. All of this has created a flurry of activity and commentary related to the understanding of human-written text. Journalists have sought clarification for the credibility of AI-generated news copies, publishers are arguing if the new models should be prohibited from scraping content from online sources, including academic paper repositories for training or making it available in a chatbot kind of interface for users. What about students who use these tools to produce essays, and do such examiners are able to identify that? Or, universities issuing bans against students accessing tools claiming it is undermining the quality of education.

Or the danger of deepfakes, where an AI text generation model is fed a huge quantity of well-known personalities, speaking on social or political issues, and creating personalized and fabricated text that comes to audiences as prevalent on behalf of those personalities. Many of these models are being launched by startup companies and large

IT firms as well. Writing assistance is also becoming more available with integrating such text writing suggestions into tools.

4.4.2. AI-Assisted Design

Generative AI is anonymously creating content, such as art, designs, games, film, and more, included in many key areas of creative work. And all of it is really an AI-assisted process. Sure, it's able to produce unique images or music off a prompt but the world is already filled with these, but the thing that it enables is the creative brainstorming process. In early or mid-phases of creative projects, say character design for a video game, an artist can prompt an AI model with a couple of key pointers on the visual style, themes, color palette, types of faces, etc., and get a row of interesting and unique designs. These can then be post-processed, fitted, and touched-up further since an AI-in-the-box can not yet generate ready-to-use and flawless works. For projects that are more template-driven, such as can be used to generate thousands of ready-to-use designs for motion graphics, ads, and promo campaigns. These have already exploded around the world, and it's very likely that businesses are ready-booting.

Applications in UI and UX design are also surfacing with platforms which makes the generation and prototyping of low-fidelity applications much easier and faster. Research has also dived into VR and 3D mods with tools. Although the process in this space is quite early and not very polished yet, it still has promise and starry-eyed fundamentals, enabling people to jump into creative 3D assets much easier and faster compared to the existing process loops. Each of these platforms has enabled businesses to bootstrap creative projects that would have otherwise taken hundreds of hours of collaboration, iteration, and fine-tuning, like prototyping an unconventional mobile app, creating ads and promo campaigns to massively accommodate the distinctiveness of every local market, generating unique assets that would have otherwise needed a hiring and collaboration cycle of weeks and months, etc.

4.4.3. Music and Art Generation

Artwork sits at the intersection of intellect and emotion. The generation of two-dimensional and three-dimensional art using Generative AI offers a unique perspective on the relationship. Generative Adversarial Networks first emerged about a decade ago to create two-dimensional images and sparked interest in various models. Now, models can generate photorealistic artwork from a text prompt. Recent diffusion models make even higher quality images. Text-to-image artwork creation can be accomplished via either of two methods. Users can either use open-source code and sample images to run the model locally or via the model publisher’s cloud service.

Three-dimensional art creation suffers from more limited operations, leading to less high-quality results. Current options include various models with text prompts. Artists train these models using a curated selection of referred artwork images, and they use either a dataset of user-uploaded real-world images or a selection of images of selected referenced artwork. Music generation has a longer history with dedicated entities producing original pieces. Some use GAN technology for seemingly infinite stream-based music production. Generative AI in music is especially suited for sales and social media jingles, video background music, or custom wedding music. Companies can lease music or proprietary music models can be used instead.

AI-created music does not currently seem likely to gain traction on top billing. Decisions about how to begin and end a song or the feel either increase in complexity with song progression, easing, or change with connecting notes, ambient sound, and accompanying instruments currently seem outside the capability of Generative AI models. Wrapping of existing song components into new samples and mixing well-conceptualized variations alongside existing artist songs appear more viable as near-term uses.

4.5. Transforming Legal Practices with AI

The legal sector is often excluded from discussions regarding technological advancements in other areas. Lawyers often resist changing their practices, believing that doing things the way they have always done them is the best way to mitigate risk. However, generative AI technology is gaining traction in law, being heralded for its potential to revolutionize access to justice and the delivery of legal services. As most other sectors, the legal industry has also been undergoing a massive crisis which has exacerbated the problems of sky-high hourly fees, inefficiency, and lack of transparency; shortages of practitioners in specific areas, such as immigration law; lack of representation in others, such as criminal law; love of other models or delivery, usually aiming for higher profits with fewer concerns about the end-user, thus creating big tech monopolies.

Document review is a laborious and, to some extent at least, a low-risk activity traditionally done by junior associates, often involving thousands of documents that need to be read to extract privileged material. The task is, however, often also implicated in firm burnout, as associates gruel over cases with effects on their work-life balance. AI tools have been designed to accommodate the need for speed, now promising to turn what took weeks to accomplish into a matter of a few hours, a boon to both firms and their clients. Updated models of generative AI are now capable of answering questions about actual legal texts or analyzing pieces thereof before a human has had a chance to look at them, being fine-tuned specifically for this kind of research. Tools specializing

in the analysis of case law have appeared in the legal landscape, some of them already being used by big names in large corporate litigation.

4.5.1. Document Review Automation

In legal practice, a variety of tasks require humans to read lengthy documents for relevant information. Document review is a time-consuming, tedious, and expensive step in litigation. Copies of emails, texts, memos, and other materials generated during prelitigation negotiations and interactions often total hundreds of thousands of pages. It is not unusual to find tens of millions of pages of documents relevant to litigation. Historically, document review has been performed by attorneys and paralegals who read, consider, and filter each document. More recently, in-house electronic discovery teams have often employed contract discovery vendors to perform the task. These vendor teams still rely heavily on human reviewers but also take advantage of machine learning algorithms specifically designed for document review. AI algorithms base their selections on the work of the initial reviewers who evaluate simple coding rules to develop test sets for training the algorithms. Usually, the AI algorithms are used to assist skilled legal personnel working with teams of contract attorney reviewers working around the clock doing speedy review of pages of discovery. Whether human-assisted or fully automated, AI-driven contract document review tools are faster and less expensive than traditional document review.

Lawyers in practice need only to learn how to use the discovery tools properly. The AI algorithms for document review have been in existence for more than a decade and are now well-established processes. Both federal and state rules of evidence provide for discovery of relevant information, and recent developments give parties codified guidelines for addressing discovery issues in litigation. Some companies specializing in AI-powered contract review tools are all of which state that their algorithms provide faster, cheaper, and more complete results, allowing lawyers to spend their time on more valuable tasks.

4.5.2. Legal Research Improvements

Advancements in large language U-models have also enabled enhancements in legal research applications. Extensive improvements in research tools allow for much seamless and intuitive searching of primary and secondary legal authority based on natural language queries and offering more precise results. Natural language processing systems increase the ability of AI to determine the meaning of a query and respond with relevant results that fit the user's intention, often offering supplemental content that may not specifically match the search query. Existing legal research tools have limited

abilities to rely on a user's input statement and have been hampered by difficulties in interpretation of nuances and context. Query-based searching often fails to return all the authoritative content relating to a legal issue. Natural language processing power is now capable of directly usable AI-enhanced research solutions for routine legal research publications. General ideals involving how legal problems or questions are framed in natural language can now enable the retrieval of relevant authority from available databases that include all substantive legal text, inclusive of statutes, case laws, agency regulations and decisions, and even secondary legal references like dedicated treatises and opinion pieces examining the issues.

4.5.3. Contract Generation and Analysis

All contracts have certain commonalities, so it is possible for AI to understand which kinds of clauses are important to a person for different contract types, and as such, tailor the content to fit that person's needs. AI models can produce simple to complex first drafts of contracts in different styles. These tools can produce work that has significant passable quality for small contracts, something like a simple NDA draft. However, lawyers should be cautious in using them and use the output with extreme caution, as the cost of mistakes or omissions in contracts is quite expensive. Nonetheless, just as translation tools are not good enough for professional translators, AI may be a tool for professional contract drafters, to help streamline their workflow.

Generative AI and NLP can also analyze contracts and help lawyers pinpoint changes and revisions. Like other generative AI products, those based on NLP will be better than what was offered before at scale as long as the user employs best practices and provides the proper initial input. Contract drafters will still want to be conservative in the generation of their documents, and will probably wish to trust but verify when working with AI's outputs for those more miscellaneous contracts. As AI has shown increasing capability for drafting contracts, it follows that it may also help with the other important task for contracts; contract analysis, with these tools significantly accelerating the time it takes to sift through documents.

4.6. Ethical Considerations of AI in Business and Law

The rhetoric about the woeful lack of ethics in business and law has reached a crescendo in recent years, and it remains to be seen if that will translate into any meaningful change. The legal environment has shifted. Corporations are more receptive to the concept of corporate values beyond simply maximizing profits. The corporations increasingly work with plaintiff lawsuits that threaten their reputational capital and may resonate with

consumers. AI is permeating business and the legal profession in so many ways that there will be hot and cold spots for regulation and advocacy.

For so-called Business Intelligence, with an emphasis on data analytics, there is a more limited scope for ethical considerations compared with Creative Processes or Law and AI. The projects should be architected and executed with an eye toward ethical behavior, but the legal penalties for low-quality work or misuse are clearer and stronger. Ethical and legal infractions related to privacy, proprietary data and intellectual property rights, trade secrets, regression to the mean or further "dumbifying" dumb normalized data, and other areas punish bad actors or offer recourse to the victims. Large project budgets, more limited review and audit pathways, and potentially more serious consequences for decision-making bodies not understanding what the computer tracking and predicting is basing its work relative to, make it prudent to have a transparent process, detailing inputs, algorithms and equations, expected outputs, and provisioning paths for reconsideration.

4.6.1. Bias and Fairness

Artificial Intelligence in Business and Law has augmented the relevance of fairness, particularly through bias - a critical factor associated with ethics. Bias, originally an effect at the human level impacting our mind, environment, and behaviors, is now an effect at AI that triggers decisions prompting impact at human level. In particular, bias is applied to a specific system - AI systems are said to be biased if they trigger decisions that do not guarantee fairness. With the business and legal applications of AI, fairness is a matter of life or death.

When talking about bias, three distinct phases may be perceived. At the first phase, bias in AIBL was ignored - there were little to none risk associated with it - it was believed that AI was 'doing the right things' from the ethics perspective. During that moment, AI was less questioned, destined to be focal for a large number of adoptions - in particular, AI was perceived to contribute to safer decisions. People tended to take confidence in AI decisions - at that moment, AIBL was considered to perform better than traditional rule-based decision-making processes. At the second phase, what we could call proof of concept phase, several AIBL real-world applications confirmed it was unable to perform better than traditional rule-based decision-making processes - sometimes with a far worse impact in the business or legal world. The impact of the adoption of AIBL systems did not always correlate with the expected value usually associated with the adoption of AI. Moreover, the behavior of AIBL models did not usually comply with our expectations. During that phase, AIBL systems started being questioned in terms of their capability of delivering ethical and moral output values, as they were perceived as an amplification of human bias.

4.6.2. Transparency and Accountability

Transparency and accountability are considered important safeguards when deploying AI inference chains, particularly within the commercial and legal domains, because the models have been described as offering results that have aspects best described as "magic" and "dirty magic," respectively. While some degree of arcane knowledge may be beyond the skill level of end-product consumers and even field experts, some product documentation and discussion would permit other professionals and some domain experts to have some understanding about what the various models are doing and whether and how they are performing relative to other alternatives. Supplying audit logs of the model decisions may be necessary. This type of transparency is especially important for models primarily offering task-specific solutions that are not adaptable for other prospective tasks. This aspect of transparency is also relevant to how adaptable the models are to corrective input that makes themselves better targets for transparency through logging or other explanations. Humans are also a vital aspect of AI systems, often providing the final quality assurance gates. Therefore, requiring that a human-in-the-loop be present in critical decision paths will help with transparency and trust. The human and AI decision-makers should be collaborating to ensure that the pair's decisions can come together to assure the best outcomes possible. This may entail training effort both on each firm's decision-makers as well as educating the AI model through cadre feedback functions to generate outputs that best align with the firm's values. Accountabilities for both the human and AI elements in the final decision cannot and should not be avoided. While the AI providing the results or even being an important partner in the resulting judgments may take some burden off the ultimate responsibility from the human, it should not be forgotten that neither should be free of accountability.

4.7. Challenges in Implementing Generative AI

The implementation of generative AI systems, and especially those for business intelligence and creativity, is fraught with challenges. One strong reason for these challenges is the relative incompleteness of the technical building blocks. In particular, while large language models, as well as foundation models in vision and audio have transformed their respective domains, we are far away from a universal model which can be adapted to every possible application. Consequently, for every new application, there are hurdles related to the type and scale of data required for fine-tuning, as well as access to the right expertise in prompt engineering, training, data labeling, and especially data curation. In addition to large language models, popular tools for generative AI include text-to-image models, the various Deepfakes, text-to-object, 3D, and video models, and music generation tools. While these have been successfully leveraged for structures or semi-structured formats of business intelligence and creativity, the domain remains

challenging for generic models beyond large language models. Similarly, while integrated ecosystems allow cross-fertilization between the available tools, generative AI ecosystems with tools, databases, user interfaces, and processes that match the productivity and efficiency of existing software are necessary before generative AI technology makes a significant dent in traditional tools for business intelligence and creativity.

An often overlooked reason for the lack of adoption is the resistance of people in the affected areas to change. For example, while a generative AI system may solve a knowledge or language-based creativity task, resistance to the tool exists due to lack of trust in the output quality, lack of an interface to customize the output to an individual style, desire for ownership and copyright of the output, and desire to continue the creative process, rather than having a ready-made product.

4.7.1. Technical Barriers

Building a high-performance application on top of a powerful LLM or generative model is not yet trivial. It requires significant Machine Learning (ML) and software engineering expertise. The cost of deploying and operating a large model serving infrastructure remains high. For fine-tuning of LLMs, amidst quota restrictions by LLM providers, the investment is carefully done and is more suitable for larger companies. The challenges of manual prompting and lack of control over hallucinations can lead to compromised ML applications not meeting certain production quality standards set by companies. The ability to let the models create new content raises important issues of governance, security, and risk, requiring new technical solutions. How does one block unallowed or wrong use of the systems, especially large publicly available models? A long tradition of work in AI on designing expert systems or on AI architectures that allow user input and interaction for particular tasks is now more important than before. This line of research helps tailor generative AI solutions to production use.

Moreover, performance on certain tasks and even particular inputs is still disappointing for large models. Many traditional ML techniques are shared to boost performance and make generative models more accessible. Such improvements in performance, trustworthiness, and use reduce the hurdles in using generative AI, complementing traditional tasks in analytics and ML approaches. Traditional rule-based programming tools using declarative specifications to better communicate and make explicit the design choices can be combined with generative AI systems. One of the challenges a programming language is a multi-shot interaction for translating domain expertise in a particular area to pre-structured code in that domain. Integrating Generative AI in commercial software development facilitates tailored services. As more APIs become

available for building Generative AI based services and applications, managing those services and tools becomes important.

4.7.2. Resistance to Change

Organizational inertia refers to the ingrained tendencies of established firms to maintain existing strategies, structures, processes, and systems. This inertia becomes a hindrance whenever changes are warranted, and is strong in response to disruptive innovation. Large firms are particularly prone to experiencing this inertia, stemming from physical and technological dependencies that take weeks or months to overcome, and resource capital dependencies wherein senior leadership often has considerable wealth at stake, leading to "calculative" forms of resistance. This resistance is fueled by the status quo bias. Indeed, decades of research on decision-making confirm that those who oppose disruption will act to preserve existing and often ineffective technologies and ways of working.

While this inertia lightens the impact of all types of disruptive market innovation on established firms, inertia applied to organizational work processes can hinder and perhaps nullify the competitive advantages offered by a firm's novelty, efficiency, and resources. As a practical matter, most managers do not have the luxury of observing established firms become disrupted and decision-making researchers have found that many practical decision-making challenges depend on underlying organizational inertia stemming from the above causes. More research on the other influences on strategy that redirect management's focus more or less automatically toward competitive positions to mitigate is needed. The linkages from one to the other should be core to their managerial challenges, including insight buyers, performance measure systems, expense budgets, overall performance comparisons, and plant tour understandings.

Heeding this latter warning, the special problem of radical innovation is a challenge for the owner-CEO entrepreneur who has an exceptionally high tolerance for variance-not-explained market uncertainties. We reemphasize the central concern for executives: a mandate to create, stimulate, and monitor not just incremental, but radical innovations, driven solely by the outside environment changes, to maintain the positions of organizational success across their firm bodies. But, perhaps mercifully, evolutionary pressure fluctuates through time.

4.8. Future Trends in Generative AI

Recent data suggests that the generative AI market is expected to make significant progress over the next few years. It was projected to develop from USD 36.5 billion in

2023 to around USD 110 billion by 2030, at a compound annual growth rate of 16.2% over the period. Contributing to this increase are a strong demand for generative AI across various industries, improvements in computing power, along with the rapid progress in advanced neural models.

However, as generative AI nears this transformative state, various concerns have arisen about its deployment, most stemming from its classification as a frontier model, its potential for misuse, and the threat it poses to many industries. Companies have been prompted to declare moratoriums on their training or release. This rising disquiet has led governments to move towards the proactive regulation of the field, stopping the activity of rogue developers while at the same time, protecting it from undue slowdown. In line with its commitment to developing safe and responsible AI, the US government signed an Executive Order to support policy development in a number of critical areas. The Order makes provisions for requirements, emphasis on collaboration with developers, and the establishment of processes to ensure safety.

Future industry developments can also be expected to evolve from the key players steering the development of the field, with the proliferation of other large companies and start-ups – notably in the West, other countries, and non-English languages or functions – in addition to the generation of models as the envisaged model landscape. Although such developments are primarily technological, they present multi-disciplinary issues in respect of their user communities.

4.8.1. Advancements in Technology

Recent advancements in foundational AI models have renewed interest in generative approaches to many AI tasks. Impressive high-quality generation of natural language text has been observed. Other foundational models include language models for dialogue applications and a generative adversarial model on enormous datasets that can synthesize high-quality images of objects that look like real photographs. A model uses generative answering to improve question-answering systems. Generating images from textual descriptions is a hard task that traditionally required artistic knowledge and talent. A method estimates the possible success of various generative processes and selects the most promising to focus on. The AI research area called "Self-Supervised Learning" aims to improve the quality of generative training by using image data to predict the next images in the data. The foundations of generative approaches take inspiration from the field of physics.

Various trends point toward accelerated exponential advancements in generative technology, driving investment and leading to both new applications and new risks. Hardware accelerators and increasing access to computing resources via cloud-based

neural networks are driving the growth of computational and optical capabilities for both vision and language tasks. Evolutionary neural architecture search approaches are continually discovering faster and smarter neural networks. Transfer learning enables expert neural networks to be quickly deployed for novel data domains and novel types of reasoning beyond what they were originally trained for. Meta-learning capabilities predict the success of models and model architectures for various datasets and tasks. Strategic competition between major governments and among major companies and among many universities is increasing resources devoted to AI capabilities.

4.8.2. Evolving Regulatory Landscape

Traditionally, laws such as copyright law have protected only works that involve a minimal amount of human creativity. Since AI is not a physical entity, it lacks the capacity to present creative elements to the work and, thus, has no rights. The problem posed by AI is that it lacks the soul that makes us human. It can replicate human actions through the mechanization of human thought, and this could lead to the disappearance of human beings with original ideas. Therefore, there is consensus in that many aspects of AI tools require regulation to protect the artist or creator, such as their trainings, to protect the authorship of its prompts, and to regulate the harms to society posed by its capacity to create content. This regulation can be done by regulating the product, the AI model, or the AI provider.

Currently, several bills are being reviewed in different parliaments and congresses. In the United States, a proposed model aimed particularly to regulate big tech stands out. It proposes a set of measures that ensure that big tech companies that benefit from having used public content for training, mandate them a tax that funds technologies to detect AI-generated content and to educate artists about these technologies. Substantially similar proposals have also been presented. Also, open letters have been signed pushing for an AI-specific regulation. With the rush towards AGI, it calls for a six-month moratorium for research to provide the opportunity to create rules in AI systems more harmonized with the values of society.

4.9. Case Studies of Successful Implementations

Through this work, we have suggested that generative AI is a powerful technology that supports the automation of both analytic and creative activities in the three sectors of business intelligence, creative processes, and law. In this section, we illustrate these functions of generative artificial intelligence and its impact on areas of work through a series of case studies of significant implementations. These cases illustrate the positive

and negative effects of generative AI in these three sectors of business, and they provide examples of how to manage the technology appropriately.

The goal is for the reader to better understand the potential for disruption as well as the real efficiency gains in work product through use of this type of AI. In the business intelligence case, we juxtapose the quantitative result of the AI Disruption Index with an accounting of work losses for financial analysts to contrast the lack of up-to-date, concrete loss data with potentially more forward looking and predictive data of the Disruption Index. In the creative process case, we examine a Ghostwriter example related to NFL coverage and how the implementation of AI to aid writing leads to a general discussion of the ramifications of generative AI-assisted first draft writing in a rewards economy. More broadly, we take a step back and try to assess what "good work" should look like in a generative future. Finally, in the legal industry case, we present the application of technology within the popular research system as a more defined instrument of work efficiency.

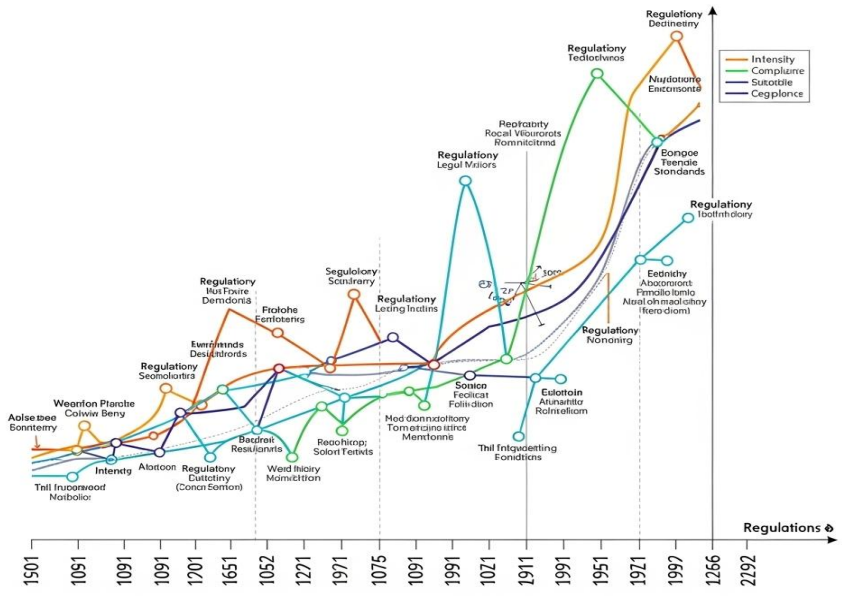

Fig : Evolving Regulatory Landscape

4.9.1. Business Intelligence Case Study

Generative AI's ability in transforming data into simple human-digestible language, significantly elevating users' productivity, is perhaps easiest to demonstrate in the area of business intelligence. Most organizations have sizeable data lakes of transactional, operational, and reporting data. Business manager users know that they should use these data to help drive significant improvement in their organizations. However, increasingly sophisticated analytical tools have been focused on data specialists having advanced technical skills — data engineers and quant analysts — who develop reporting dashboards consumed by the business manager users. These dashboards have become the lingua franca for BI but have significant limitations: 1) typically very few; 2) each dashboard trying to provide answers to a few questions only; and 3) the answers are displayed in static look-ups or visualizations that the manager must interpret. This problem is exacerbated by the focus on large data needing specialist tools.

Over the last few years, Natural Language Processing, a sub-field of AI that has made significant progress, has been duly incorporated into existing BI-specific applications. However, large generative models are rapidly changing the landscape. They power both open-source models and SaaS BI applications that provide an entirely new user experience that uses AI's capabilities to address BI's limitations. Already, we see organizations formerly squelched by the dareness of using internal data beginning to dip their toes in the water.

4.9.2. Creative Industry Case Study

While the law sector braced for disruption from generative AI tools focused on contract analysis, brief writing, and legal research, the traditional creative economy, in advertising and marketing, was first hit with the generative AI wave with tools focused on general and commercial image generation. These tools allowed marketers, designers, and artists to effortlessly and cheaply generate unique and interesting images to use in communications but at the same time risked drowning them within a tide of subpar generative images flooding attention elastic social platforms.

Needless to say many marketers and designers quickly adopted the tools to create compelling and often unique visuals with the output, in many cases, rivaling skilled artists. The ease of use and speed of production allured advertisers and marketers who could harness the new tools. But as many agencies and brands jumped on the bandwagon, the oversaturation of weak, vapid, and sometimes grotesque AI-generated visuals soon turned rules to the creative economy. The marketplace balance soon shifted with demands for the original distinct applied art and craftsmanship of real artists whose output came to represent a brand’s true unique identity, sense of purpose, and values.

Creative commissions and consulting hustles catered to discovering that sharp edge. However, demand was still unbelievably strong for production assistance in formulating quick ideas as well as budgetarily-friendly two-dimensional and three-dimensional visuals.

4.9.3. Legal Sector Case Study

Legal professionals are today offered generative AI systems that let them "prompt" the system with a short instruction that encapsulates everything they need to outcome. An example of this type of prompt could be: "I have a case where the suing party claims damages arising from a defective product. I note the descriptions of the defective product in the complaint of the suing party. As an assistant, I would like you to help me by comparing the description of the defective product with the description of other similar products that I have noted underneath this prompt. After you analyze the comparison made, please elaborate on a description of the similar products as independent operators to a normal operator of products of that category. From then, please tell me if you consider that there are enough elements – based on the similarity of the description of the defective product and the description of the other products that you have been analyzing under this prompt – to define that the products that you are analyzing are indeed associated as competitors with the defective product and, in consequence, merit the same responsibilities of the suing party described in his/her complaint." The generative AI processes this request, analyzes it, and then proceeds to produce what are – or at least could be considered – accurate legal analyses.

Clearly, this is a repetitiveness type of task, and the opportunity to ensure that it is sufficiently produced at the sector level could help empower lawyers to bring more value to their clients and society. Now, while intuitively, this is a low conjecture regulatory task that could be easily mechanized, this view must be considered by regulatory frameworks duly designed to inhibit the relevant risks that this task entails in a sector as sensitive to society, economy, and its development as is the legal sector.

4.10. Conclusion

Supporters and detractors alike question whether generative achievements add enough value to justify their costs. Increasingly, the answer appears to be yes. Advancing business intelligence, creative processes, and legal practice, generative AI has immeasurable potential. It accelerates time-consuming tasks, enhances creativity and productivity, and increases work quality. Its potential business use cases within business intelligence as well as in creative fields appear wide-ranging. We can clarify text; visualize, analyze, and manipulate data; and derive insights from various data sources.

Generative AI can design fantastic images quickly and with few limitations, create entire scripts or screenplays, compose riveting music or realistic sound effects, and even complete exciting, coherent, and imaginative video output. For those within the law, generative AI can summarize and research statutes and regulations, predict litigation and criminal case outcomes, draft contracts, outsource due diligence, automate reviews with high accuracy, and manage complex legal workflows.

While the traditional model of creative and intellectual output has emphasized uniqueness, distinctiveness, and originality, generative AI directly challenges this view. In democratizing creative processes, easing burdensome legal workflows, and freeing legal practitioners to focus on strategic, complex, and highly manual work, generative AI has immeasurable business potential. However, because it cannot approach humans in the areas of nuanced perspective and reasoning, real-world empathy, and original thought, it is marketers, lawyers, and other professionals who must carefully articulate the exact querying and input process, understand the tools' inherent limitations, and use the tools as an interactive and collaborative base from which to create. For any business that treads carefully, and that incorporates generative AI cautiously and pragmatically, the benefits appear poised to vastly outweigh the costs.

References

N. Bommasani et al., "On the opportunities and risks of foundation models," arXiv preprint arXiv:2108.07258, 2021.

A. K. Goel et al., "The impact of generative AI on creativity and law," IEEE Comput., vol. 56, no. 4, pp. 56–65, 2023.

S. Liu et al., "Harnessing GPT models for enterprise decision-making," Proc. IEEE Int. Conf. Big Data, pp. 3512–3521, 2022.

M. McCormack, "Generative AI in the legal domain: Current applications and ethical concerns," AI Law Rev., vol. 1, no. 2, pp. 22–36, 2023.

D. T. Liu and A. S. Jatowt, "Generative AI for business content creation," IEEE Access, vol. 11, pp. 45126–45136, 2023.