

Explainable Large Language Models & iContracts

Georgios Stathis^a

Institute of Tax Law and Economics, Leiden University, Kamerlingh Onnes Building,

Keywords: LegalTech, Intelligent Contracts, Large Language Models, Explainable AI, User Trustworthiness.

Abstract: Contract automation is a field of LegalTech under Artificial Intelligence (AI) and Law that is currently undergoing a transition from Smart to Intelligent Contracts (iContracts). iContracts aim to full contracting automation. Their main challenge is finding a convincing direction for market adoption. Two powerful market factors are the advent of Large Language Models (LLMs) and AI Regulation. The article investigates how the two factors are able to influence the market adoption of iContracts. Our Research Question reads: *to what extent is it possible to accelerate the adoption of Intelligent Contracts with Explainable Large Language Models?* Following a literature review our research employs three methodologies: market gap analysis, case study, and application. The results show a clear way for iContracts to follow, based on existing market gaps. Moreover, they validate whether the application of Explainable LLMs is possible. The discussion clarifies the main limitations with Explainable LLMs. Our conclusion is that the two factors are impactful for so long as the market adoption attempts to bridge the gap between innovators and early adopters.

1 INTRODUCTION

Smart Contracts have laid the foundation for iContracts. Smart Contracts are self-executing contracts with the terms of the agreement directly written into code (Madir, 2020). They are rooted in blockchain technology, which provides transparency and trust in the digital realm (Werbach, 2018).

While Smart Contracts have revolutionised the contracting process by automating transactions and reducing the need for intermediaries, they have limitations. Smart Contracts are essentially binary, and capable of executing predefined actions, but incapable of interpreting and adapting to complex legal nuances (Mik, 2017).


iContracts represent the next step in the evolutionary process. They are designed to go beyond the straightforward nature of Smart Contracts. The key objective is to enable iContracts to handle autonomously the entire contracting process, encompassing everything from negotiation to execution. Such an action path entails a seamless integration of human-readable language and code (Mason, 2017). As a result, iContracts will possess the capacity to understand, adapt, and evolve in response to the intricacies of real-world contracts, thus paving the way for

a new era of automation (Stathis et al., 2023c; Stathis et al., 2023b).

- **Definition 1:** An **intelligent contract** or **iContract** is a contract that is fully executable without human intervention ¹.

Meanwhile, the main challenge with iContracts is the market adoption rate (McNamara and Sepasgozar, 2018). Automation is closely related with the law. Owing to the high legal consequences of iContracts, human users prefer traditional methods such as the direct inclusion of a legal expert (Stathis et al., 2023a). Such a preference is closely connected with two developing trends: (1) the advent of Large Language Models (LLMs) and (2) the regulation of AI. The adoption of LLMs has happened at the speed of light. As a result we may wonder how long the adoption of iContracts will take. Is it five years or only half a year? Next to this question we are facing (1) the regulation of AI, with the AI Act being currently in preparation (European-Parliament, 2023), and (2) the global battle to regulate technology (Bradford, 2023). One of the main challenges with AI technologies is the lack of user *trustworthiness* (Liang et al., 2022). Research shows that by understanding how AI takes

¹<https://bravenewcoin.com/insights/pamela-morgan-at-bitcoin-south-innovating-legal-systems-through-blockchain-technology>

^a  <https://orcid.org/0000-0002-4680-9089>

decisions the outcome of the decision can then be explained and hence user trustworthiness increases (Xu et al., 2019).

- **Definition 2: Large language models (LLMs)** are a category of foundation models trained on immense amounts of data making them capable of understanding and generating natural language and other types of content to perform a wide range of tasks ².

Our motivation for performing this research is to improve the rate of market adoption of iContracts. We believe that by examining to what extent it is possible to apply Explainable LLMs on iContracts, we are able to increase user trustworthiness among humans.

The aforementioned thoughts lead us to the following Research Question (RQ).

RQ: To What Extent Is It Possible to Accelerate the Adoption of Intelligent Contracts With Explainable Large Language Models?

The article's contribution is that it shows and validates whether the application of Explainable LLMs on iContracts is possible to direct the future attempts in developing iContracts fit for market adoption.

To answer the RQ, we structured the paper as follows. In Section 2, we describe the relevant literature. Section 3 presents the methodology and Section 4 the results. Section 5 discusses the results and focusses on theoretical and practical parameters. Finally, Section 6 answers the RQ and provides the conclusion together with a preview on the future.

2 LITERATURE

The literature present sources on: the three contracting revolutions (2.1), market adoption in business studies (2.2), the application of LLMs in contracts (2.3) and the development of trustworthy and Explainable LLMs (2.4).

2.1 Three Revolutions

The path from human contracts to iContracts is characterised by three revolutions. The *first* revolution started with the transition from physical contracts to digital contracts (also known as Electronic Contracts or eContracts) (Krishna and Karlapalem, 2008). With the advance of big data, people gradually digitised physical contracts into electronically accessible documents (von Westphalen, 2017). Moreover, they replaced some physical labor such as signing, by electronic handling. Today, most market developments

²<https://www.ibm.com/topics/large-language-models>

are concentrating on further adopting and expanding the adoption of digital contracts ³.

In the last decade (2010-2020), with the advent of blockchain technology the *second* revolution occurred, which switches the focus from digital contracts to smart contracts. Smart contracts are agreements that are executable by code, most often on blockchain-based distributed ledgers (Khan et al., 2021). The promise of smart contracts is to be seen as the replacement of human language by human code. However, the adoption of smart contracts, beyond specialised blockchain or experimental market circles, has not managed to reach automatically wide market adoption ⁴. One of the main challenges of smart contracts is that it is hard for users to understand and trust the computer code behind them (Zheng et al., 2020).

Here we arrive at the *third* revolution, the transition from smart contract to iContracts. iContracts aim at full contracting automation with minimal to no human involvement (Stathis et al., 2023c; Stathis et al., 2023b). iContracts promise to bridge the gap that smart contracts have been unable to fill. "Bridging" will gain user trustworthiness by combining computer code with user-friendly, readable and understandable code, that originates from physical contracts (McNamara and Sepasgozar, 2020). Despite the surge in scientific interest in iContracts for the construction industry, there remains a significant demand for traditional contracting (first recognised in 2020 (McNamara, 2020)); meaning the development of iContracts is still in its formative stages.

As a follow-up, we mention a research proposal (Stathis et al., 2023c; Stathis et al., 2023b) presented at International Conference on Agents and Artificial Intelligence (ICAART) 2023, which revived research on iContracts by re-purposing use cases, instead of delegating the task of the complex construction industry. A reference to a straightforward freelancer agreement would simplify the evident complexity in construction contracts. In particular, such freelance agreements stipulated the extent to which iContracts may contribute in specific domains. With increasing attention to the details of iContracts in simpler domains it is possible to gradually expand into more complex domains. Notwithstanding these developments, it is still unclear how to develop a path towards the market adoption, that contributes towards end-users increasingly adopting iContracts.

³<https://www.gartner.com/en/documents/3981321>

⁴<https://www.grandviewresearch.com/industry-analysis/smart-contracts-market-report>

2.2 Market Adoption in Business Studies

In business studies, market adoption has been studied for many years (Posthumus et al., 2012; Botha, 2019). The most important contribution is Geoffrey Moore's *Crossing the Chasm*. It studies how technologies penetrate a group ranging from (1) innovators to (2) early adopters and gradually advance into (3) wide adopters, (4) late adopters until they reach the so called (5) laggards (Moore, 1991). The results are general: every technology that enters market adoption follows the same route (Goldasteh et al., 2022). We are going to examine the adoption of iContracts through the same five lenses. Each contracting revolution can be located at a specific stage in Moore's *Crossing the Chasm*. By identifying the specific position of iContracts relative to digital and smart contracts, it is (1) also possible to identify relevant gaps that require to be bridged, and then (2) to investigate the extent to which LLMs can contribute in bridging the gaps.

2.3 Large Language Models in Contract Automation

Since the dissemination of LLMs with ChatGPT, scientists and business people have attempted to apply LLMs, falling under the category of Generative AI, on contracts (see for example recent eighty million fundraising for Silicon Valley based Harvey⁵). The extent and results of such applications are still largely unclear. On the one hand, end-users expect that generative AI tools assist them in developing contracts at a fraction of the costs that it would usually cost to pay a lawyer⁶. On the other hand, the generated contracts (seen as LLMs) are not necessarily perceived as *trustworthy*, as with most outputs of LLMs and Generative AI (Lenat and Marcus, 2023). They may find results vague, confusing or even mysterious (potentially the result of hallucination). Hence, due to the high impact of legal consequences, end-users (even if they do attempt themselves to generate a contract) still prefer to consult their lawyer (Stathis et al., 2023a). The *larger* the degree of the severity of legal consequences, the *lower* the reliance on generative contracts (Stathis et al., 2023a). Some scientists might argue, that in a way Generative Contracts are announcing the *fourth* contracting revolution (Williams, 2024). From our

⁵<https://siliconangle.com/2023/12/20/harvey-raises-80m-build-generative-ai-legal-professionals/>

⁶<https://www.docusign.com/blog/products/generative-ai-contracts-agreements>

perspective, we consider Generative Contracts to be only a feature (i.e., one aspect) of iContracts (or of Smart and Digital contracts) and not the architecture behind iContracts. In the methodology, result and discussion sections we will clarify why we believe that this is so.

2.4 Trustworthy and Explainable Large Language Models

In 2024, we may expect that the European Union (EU) will approve the very first *rules* for AI in the world (European-Parliament, 2023). The rules will be presented in the AI Act, a regulation which is aiming to regulate two kinds of AI: Predictive AI and Generative AI (Stathis and van den Herik, 2023). The main focus of these rules is to make sure AI can be trusted in accordance with the directions provided to the EU by the High-Level Expert Group on AI (HLEG-AI, 2019). Here, we see two sides: (1) society seems to miss trust towards AI, and (2) the EU wants to protect its people by developing rules to prevent negative consequences by the wide adoption of AI (Lockey et al., 2021). This is precisely the reason why *some* AI technologies are completely prohibited and *other* AI technologies are considered as High-Risk, which are also subject to most of the regulatory restrictions⁷. *Generative AI* are treated as an exceptional AI technology which should adhere to specific restrictions, due to the high reliance on trained data which are potentially subject to copyright laws⁸ (Helberger and Diakopoulos, 2023). Due to the difficulty in deciphering levels of *perceived* trustworthiness from *actual* trustworthiness, scientists emphasise the concepts of transparency and accountability in matters related to trustworthy AI (Munn, 2023). The main principal way along which an AI-system can gain transparency and accountability is *via explainability* (Holzinger et al., 2020). Explainable AI is the field in which researchers investigate how AI system decisions can be explained in an *understandable, interpretable and trustworthy* manner for humans⁹. In the Springer AI and Ethics Journal researchers deal with the question: how the explainability issue can be handled with

⁷<https://www.consilium.europa.eu/en/press/press-releases/2023/12/09/artificial-intelligence-act-council-and-parliament-strike-a-deal-on-the-first-worldwide-rules-for-ai/>

⁸<https://www.consilium.europa.eu/en/press/press-releases/2023/12/09/artificial-intelligence-act-council-and-parliament-strike-a-deal-on-the-first-worldwide-rules-for-ai/>

⁹<https://www.marktechpost.com/2023/03/11/understanding-explainable-ai-and-interpretable-ai/>

the development of a "culture of explainability" that supports the explicit explanation of the architecture, design, production, and implementation of decision-making across the AI development and implementation chain (Stathis and van den Herik, 2023). The result is very important, even vital for our research, provided that for the case of applying Generative AI on iContracts, such explainability culture will impact the trustworthiness of the end user significantly. As positive implications we particularly foresee positive developments for ethical and legal transparency and accountability.

3 RESEARCH METHODOLOGY

The research methodology starts with an analysis of the market gap (3.1) and then introduces the case study (3.2) as well the Generative AI application in close coherence with LLMs (3.3).

3.1 Market Gap Analysis

Our methodology begins with the visualisation of the three revolutions into a graph based on Moore's Crossing the Chasm graph (Figure 1). This visualisation will allow us to identify the positioning of the three technologies (digital, smart, and intelligent contracts) in relation to market adoption. Thereafter, we are able to conduct a gap analysis to identify the steps that appear to be the missing links for the adoption of iContracts. Our gap analysis will focus on a specific case study.

To classify market categories on contracting automation we are going to use Legalcomplex's categorisation¹⁰. Legalcomplex is the largest database on LegalTech solutions. Legalcomplex has classified contract automation solutions into the following five categories: (1) contract negotiation, (2) contract risk management, (3) contract drafting, (4) contract extraction and (5) contract management.

The classification of Legalcomplex follows the typical journey of a contracting user. Starting with negotiations, a legal expert makes an estimation of legal risk and drafts a contract. Thereafter, relevant elements of that contract can be extracted during the execution stage and/or the monitoring stage; finally the contract is being discussed on management level until its completion.

¹⁰<https://legalcomplex.com>

3.2 Case Study

The case study is presented by the Knowledge Graph developed in the ICAART 2023 research and is based on the Onassis Ontology (see Figure 2). The Knowledge Graph validates how a small agreement (between two contractors, guided or supervised by a legal expert) can be automated based on *communications* and *risk* data that are exchanged between two contracting parties (Figure 2). In the communication with the contractor, the back-end the role of a legal expert is to clarify the *scope*, *contract*, *risks* and *questions* for the contracting parties. When (1) contracting parties (2) select a scope, (3) the reply questions and (4) the risk-intelligent contract is updated by the variables from the conversation exchange (see in Figure 2 the lines: question, answer, section, contract, variable). The question then is, how can Generative AI help? Our experiment will guide the reader to an answer to this question.

3.3 Application

The application shows to what extent is it possible to develop alternative options by Explainable LLMs to facilitate the automated generation of data for the stakeholders involved in our case study (i.e., the legal expert and the contractors). Starting with the KG, we will first identify the potential locations where Generative AI and LLMs can be applied. Then we are showing the power of Google's Bard LLM to *validate* whether and to what extent it is indeed possible to develop explainable generated data that functionally assist the iContract users into completing their work. To amplify the explainability capabilities of Generative AI and LLMs, in compliance with the literature on trustworthy AI, we will use the opportunity to request the LLM model to provide higher levels of user trustworthiness with alternative explanations.

4 RESULTS

The results visualise the market gap analysis (4.1) and provide the experiment validation (4.2).

4.1 Market Gap

We start with the visualisation of the market adoption of contracting technologies in accordance with Moore's Crossing the Chasm (Figure 1) (Moore, 1991). First, we see that physical contracts are the preferred contracting method for laggards. Then, digital contracts are reaching the late market. Here,

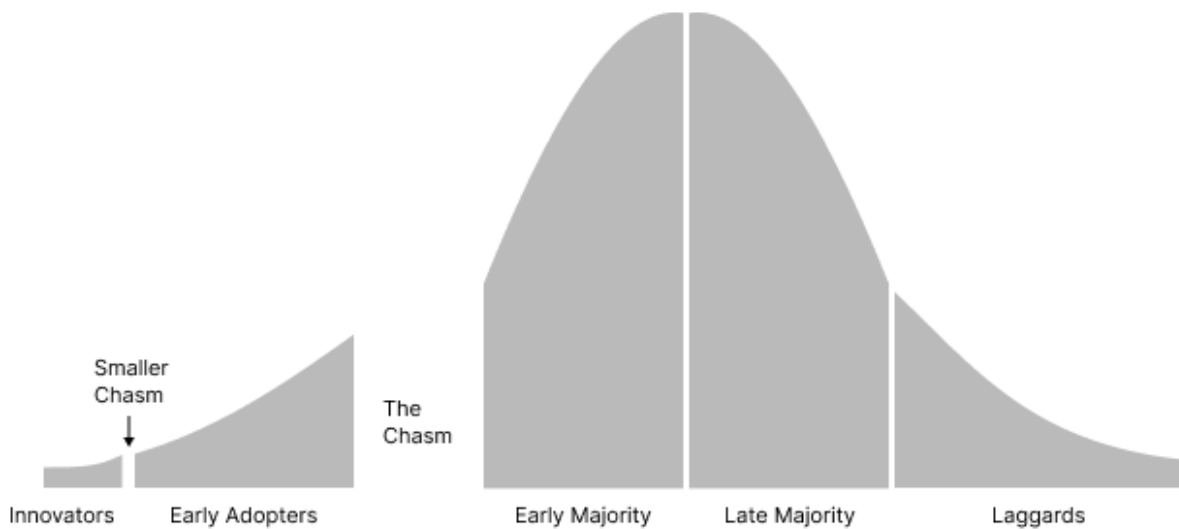


Figure 1: Moore's Crossing The Chasm.

we remark that *Smart contracts* have not managed to achieve a wide adoption of crossing the chasm. Third, at the same time, *iContracts* have still *not* reached early adopters and are therefore being evaluated by innovators only.

Based on (1) this visualisation and (2) the user preferences, we are able to identify the relevant market gap among each contracting alternative. Using as comparison framework Legalcomplex's classification of contract automation solutions, we see in Table 1 (column 2) that the highest adoption is observed with *physical* contracts. The main obstacle for physical contracts is data extraction. That obstacle has been solved by digital contracts. Hence, gradually digital contracts are the first followers. Still digital contracts have *not* managed to replace negotiations or risk management as they occur in physical contracts. Also contract drafting seems to have remained significantly reliant on physical contract drafting (although it is nowadays done often by electronic means, even with the use of templates). *Smart* contracts (see column 4) have lower levels of adoption across most categories, except with drafting, extraction and management which is amplified by electronic means. With *intelligent* contracts, we see the lower adoption rate. The table helps us to decipher which gap *iContracts* should aim to bridge first. That is the gap in contract negotiations, risk management and drafting. As seen in Figure 2, this is the conceptual line of execution that the Onassis Ontology is taking. Hence, by experimenting with the application of Explainable Generative AI and LLMs on the Onassis Ontology, it is expected to bridge the market gap observed on these three categories faster.

4.2 Explainable Generative AI Validation

The application results, which are accessible on GitHub (¹¹) validate that Generative AI can be leveraged in every single step of the Onassis Ontology. For as long as the Generative AI prompt requires the specification of supporting explanations, the provision of such explanations is possible. As a direct consequence, we see that it becomes visible that the *reduction of human labour* in *iContracts* is further amplified with the help of Generative AI. By including the explicit provision of explanations we may also expect that the user trustworthiness increases. Hence, technologically, it is possible to support the gap difference between *iContracts* with other contracting alternatives with Explainable Generative AI.

5 DISCUSSION

In our discussion we focus on the implications for the adoption of *iContracts* (5.1) and the limitations observed in the Explainable LLM application (5.2).

5.1 *iContracts* Market Adoption Implications

As we have seen, the market gap analysis became useful in identifying the precise gap between the alternative contracting options for end users. We found that

¹¹<https://github.com/onassisontology/onassisontology/blob/main/img/EGENAIEXP.png>

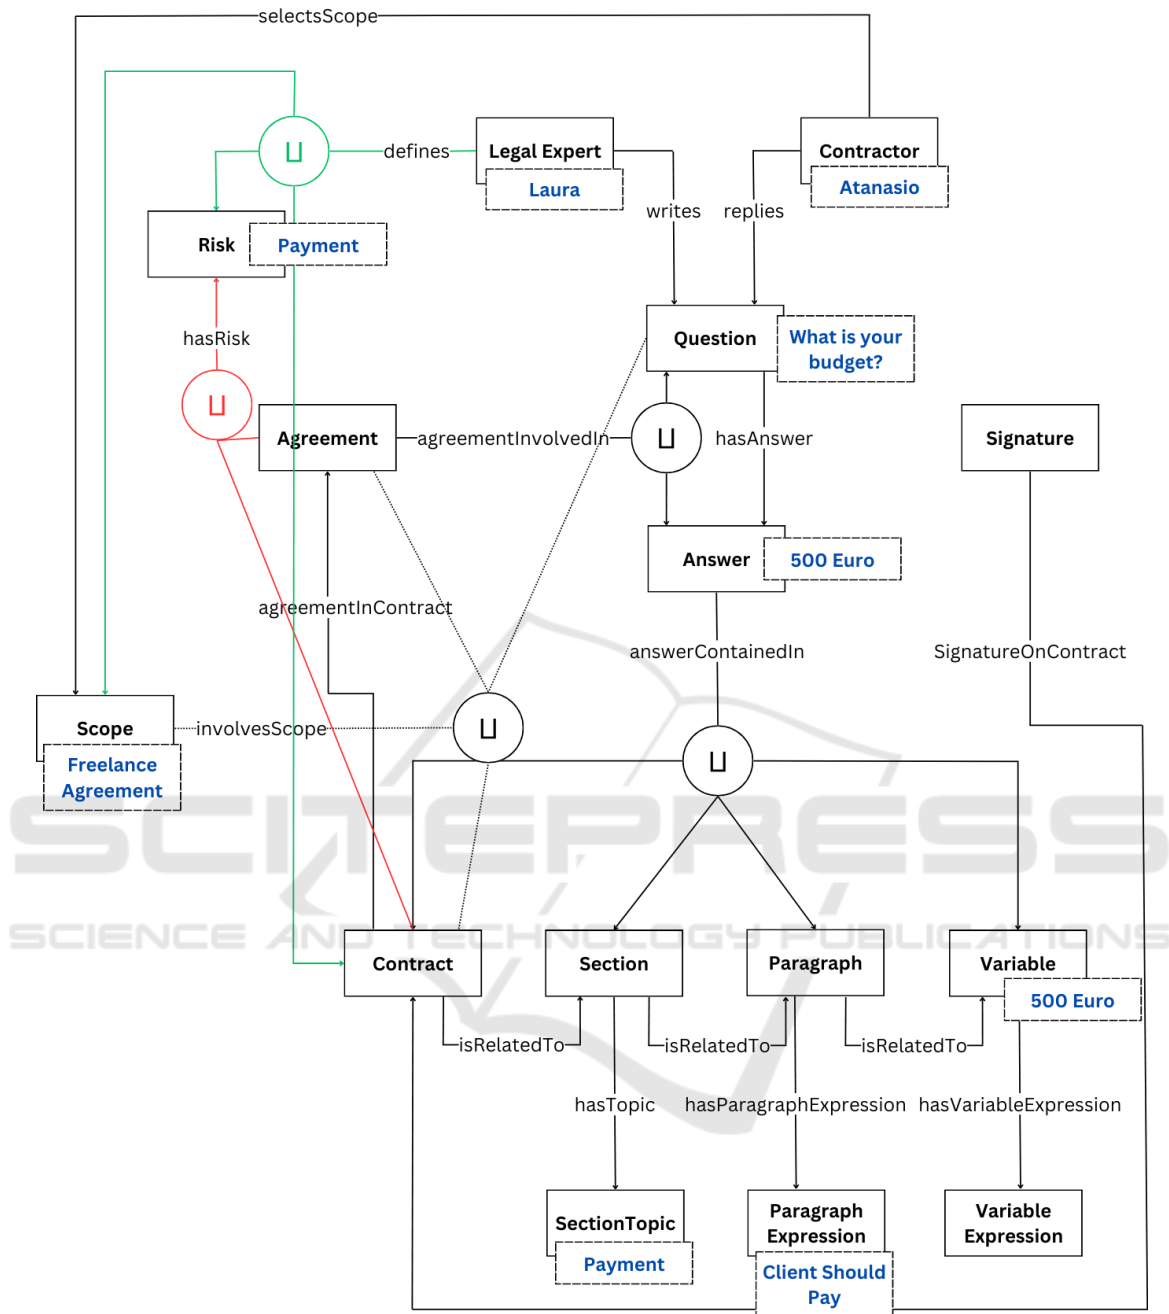


Figure 2: Visualisation of the use case scenario modelled with Onassis’s ontological expressiveness.

the advantage of iContracts lies exactly where digital and smart contracts seem to struggle with replacing Physical contracts, namely contract (1) negotiations, (2) risk management and (3) drafting.

In compliance with market adoption, business study literature, provided that physical, digital and smart contracts already present certain alternatives that users have habituated. The opportunity for iContracts to impress early adopters and gradually wide

adopters, lies in exploiting opportunities that significantly impact current practices.

In conclusion, there is thus an opportunity by applying Explainable Generative AI, which simplifies human complexity in multiple directions, as validated by the application on the case study. Still, such application allows us also to identify specific limitations which we present below.

Table 1: Rate of contracting adoption based on high-mid-low factors among four alternative contracting options.

	Physical	Digital	Smart	Intelligent
Contract Negotiation	High	Mid	Low	Low
Contract Risk Management	High	Mid	Low	Low
Contract Drafting	High	Mid	Mid	Low
Contract Extraction	Low	High	Mid	Low
Contract Management	High	High	Mid	Low

5.2 Explainable Large Language Model Limitations

Requesting an LLM to develop supporting explanations is a really possible challenge. Such explanations improve user trustworthiness by supporting their evaluation process. Instead of requesting an LLM to directly provide one single answer for a given topic, the alternative option is to allow for users to select the most serviceable answer for a specific purpose.

If we compare the effectivity of humans to produce alternative explanations against an LLM producing such alternatives, we find that it is a laborious task that most humans would avoid due to increased complexity. Hence, an Explainable LLM is able to replace a laborious task with a more trustworthy alternative.

An additional positive aspect is that we validated the application of Explainable LLM as being possible across each step of the Onassis Ontology. Indeed, we found that the application of Generative AI on iContracts happens on feature level and not on architecture level. Although, our request for explanations still presents certain limitations, which are not yet solved.

The limitations are related with the inability of Generative AI systems to provide (1) readable sources that verify its explanations and (2) understandable reasoning patterns that explain the reasoning of its developers. In compliance with the "culture of explanation" as supported in the Springer research, these two characteristics will be able to drastically increase the user trustworthiness even further, due to the ability for users to validate the origin of specific data as well as the line of reasoning a developer has followed. Such transparency helps with improving end user evaluations as well as with assigning liability with improved accountability that also increases user trustworthiness from a legal point of view.

6 CONCLUSION

The RQ of this research is: *to what extent is it possible to accelerate the adoption of Intelligent Contracts with Explainable Large Language Models?*

The answer is that Explainable LLMs, as a category of Generative AI, have the possibility to accelerate the adoption of iContracts to the extent that it applies on the categories that are currently underrepresented by the competing contracting alternatives, namely contract (1) negotiations, (2) risk management and (3) drafting. This occurs because of two characteristics for Explainable LLMs. *First*, Generative AI can make the laborious human tasks involved in the three categories simpler. *Second*, the explainability aspect can increase the end user trustworthiness significantly by supporting the outputs of Generative AI with explicit explanations. To further leverage explainability, we advice for being compliant with the "culture of explanations". Hence, the developers of LLMs should be *connaisseurs* of the culture in which the system operates when given the task to make an explicit presentation of sources supporting data outputs as well as explicit representation of the line of reasoning supporting algorithmic models.

Our further research will focus on developing experiments with additional case studies, with gradually increasing complexity. Owing to a high reliance on end user reactions, we will continue to conduct end user *validation experiments with graphical user interfaces*.

The overall novelties of the paper is that (1) it presents the state of market adoption of contracting technologies, (2) it identifies the gap among the alternative contracting technologies, (3) it connects the aspect of explainability with LLMs, and (4) it concludes that the application of Explainable LLMs is successfully possible on practical case studies.

REFERENCES

- Botha, A. P. (2019). Innovating for market adoption in the fourth industrial revolution. *South African Journal of Industrial Engineering*, 30(3):187–198.
- Bradford, A. (2023). *Digital Empires: The global battle to regulate technology*. Oxford University Press.
- European-Parliament (2023). *Artificial Intelligence Act: Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on*

- artificial intelligence (artificial intelligence act) and amending certain Union legislative acts.* Eur Lex, European Union.
- Goldasteh, P., Akbari, M., Bagheri, A., and Mobini, A. (2022). How high-tech start-ups learn to cross the market chasm? *Journal of Global Entrepreneurship Research*, 12(1):157–173.
- Helberger, N. and Diakopoulos, N. (2023). Chatgpt and the AI Act. *Internet Policy Review*, 12(1).
- HLEG-AI (2019). High-level expert group on Artificial Intelligence. *Ethics guidelines for trustworthy AI*, page 6.
- Holzinger, A., Saranti, A., Molnar, C., Biecek, P., and Samek, W. (2020). Explainable AI methods-a brief overview. In *International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers*, pages 13–38. Springer.
- Khan, S. N., Loukil, F., Ghedira-Guegan, C., Benkhelifa, E., and Bani-Hani, A. (2021). Blockchain smart contracts: Applications, challenges, and future trends. *Peer-to-peer Networking and Applications*, 14:2901–2925.
- Krishna, P. R. and Karlapalem, K. (2008). Electronic contracts. *IEEE Internet Computing*, 12(4):60–68.
- Lenat, D. and Marcus, G. (2023). Getting from generative AI to trustworthy AI: What LLMs might learn from cyc. *arXiv preprint arXiv:2308.04445*.
- Liang, W., Tadesse, G. A., Ho, D., Fei-Fei, L., Zaharia, M., Zhang, C., and Zou, J. (2022). Advances, challenges and opportunities in creating data for trustworthy AI. *Nature Machine Intelligence*, 4(8):669–677.
- Lockey, S., Gillespie, N., Holm, D., and Someh, I. A. (2021). A review of trust in Artificial Intelligence: Challenges, vulnerabilities and future directions.
- Madir, J. (2020). Smart contracts-self-executing contracts of the future? *Int'l. In-House Counsel J.*, 13:1.
- Mason, J. (2017). Intelligent contracts and the construction industry. *Journal of Legal Affairs and Dispute Resolution in Engineering and Construction*, 9(3):04517012.
- McNamara, A. (2020). Automating the chaos: Intelligent construction contracts. In *Smart Cities and Construction Technologies*. IntechOpen.
- McNamara, A. and Sepasgozar, S. (2018). Barriers and drivers of intelligent contract implementation in construction. *Management*, 143:02517006.
- McNamara, A. J. and Sepasgozar, S. M. (2020). Developing a theoretical framework for intelligent contract acceptance. *Construction innovation*, 20(3):421–445.
- Mik, E. (2017). Smart contracts: terminology, technical limitations and real world complexity. *Law, innovation and technology*, 9(2):269–300.
- Moore, G. A. (1991). *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*. HarperBusiness.
- Munn, L. (2023). The uselessness of AI ethics. *AI and Ethics*, 3(3):869–877.
- Posthumus, B., Limonard, S., Schoonhoven Van, B., and Verhagen, P. (2012). Business modelling and adoption methodologies state-of-the-art.
- Stathis, G., Biagioni, G., Trantas, A., and van den Herik, H. J. (2023a). Risk Visualisation for Trustworthy Intelligent Contracts. In *the Proceedings of the 21st International Industrial Simulation Conference (ISC), EUROSIS-ETI*, pages 53–57.
- Stathis, G., Trantas, A., Biagioni, G., Graaf, K. A. d., Adriaanse, J. A. A., and van den Herik, H. J. (2023b). Designing an Intelligent Contract with Communications and Risk Data. *Springer Nature: Recent Trends on Agents and Artificial Intelligence (Submitted)*.
- Stathis, G., Trantas, A., Biagioni, G., van den Herik, H. J., Custers, B., Daniele, L., and Katsigiannis, T. (2023c). Towards a Foundation for Intelligent Contracts. In *the Proceedings of the 15th International Conference on Agents and Artificial Intelligence (ICAART)*.
- Stathis, G. and van den Herik, J. (2023). *Ethical & Preventive Legal Technology*. Springer Ethics & AI Journal, Cham, Switzerland.
- von Westphalen, F. G. (2017). Contracts with big data: The end of the traditional contract concept? In *Trading Data in the Digital Economy: Legal Concepts and Tools*, pages 245–270. Nomos Verlagsgesellschaft mbH & Co. KG.
- Werbach, K. (2018). *The blockchain and the new architecture of trust*. Mit Press.
- Williams, S. (2024). Generative Contracts. *Arizona State Law Journal, Forthcoming*.
- Xu, F., Uszkoreit, H., Du, Y., Fan, W., Zhao, D., and Zhu, J. (2019). Explainable AI: A brief survey on history, research areas, approaches and challenges. In *Natural Language Processing and Chinese Computing: 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9–14, 2019, Proceedings, Part II 8*, pages 563–574. Springer.
- Zheng, Z., Xie, S., Dai, H.-N., Chen, W., Chen, X., Weng, J., and Imran, M. (2020). An overview on smart contracts: Challenges, advances and platforms. *Future Generation Computer Systems*, 105:475–491.