

Contracting for AI: Promise and Perils

Mari Sako

January 2023

“Modern economies are held together by innumerable contracts” which underpin high-performing and trusted trading relationships.ⁱ While data scientists, with contracting professionals, are busy developing artificial intelligence (AI) tools for contract management, not enough attention is paid to resolving an important issue, namely new challenges posed when contracting for the use of AI tools and data. This column argues that such consideration is essential for enhancing the competitive advantage of providers and users of AI tools and would contribute to public good.

We begin with discussing efforts made to apply AI and machine learning (ML) for contract management. We then shift our attention to challenges of contracting for AI by addressing a question central to this column: what characteristics of AI/ML make it distinctively difficult to contract? And what solutions exist to deal with these challenges?

Promise of AI for contract management

Contracts contain commercial terms, pricing, obligations, incentives, risk details, liabilities, and other attributes. They are drafted, negotiated, signed, reviewed, and renewed. A typical Fortune 500 company might have over a million live contracts at any point in time. For example, in 2019, Microsoft managed 1.1 million contracts in sales, licensing, non-disclosure, financing, etc.ⁱⁱ In this context, the application of AI, and ML and natural language processing (NLP) in particular, offers the promise of streamlining data extraction, improving contract oversight, and enhancing vendor compliance.

AI also promises to transform static vendor contracts into dynamic assets. AI tools may enable firms to proactively identify opportunities for new sales, improve risk management, and anticipate and prevent disputes. Combining AI with smart contracts take this idea further, with programs stored on a blockchain that run when predetermined conditions are met. The execution of agreements, for instance to release funds to specified parties, is automated with the transactional outcome revealed to all participants instantaneously.

Perils in contracting for AI/ML

Despite these promises in applying AI for contracting, when we turn to contracting for the use of AI, there are problems. In particular, contracting for ML poses extra perils due to lack of clarity over data ownership, allocation of responsibility for training the model, and commercial sensitivity undermining data aggregation. These factors impede access to good quality data, essential for achieving data-driven innovation.ⁱⁱⁱ However good the technical system, ML algorithms in themselves would not create value without training data as complementary assets. We consider below three perils in relation to data.

Data ownership

First, AI algorithms are normally kept as trade secrets. Most commonly, the vendor owns the intellectual property in the AI algorithm, whilst the customer has a licence to use the AI software. But when it comes to data to train the AI model, ownership of such data is unclear.

A core reason why ownership ambiguity exists lies in the nature of data. Unlike a physical good, data can be reused, repackaged, and resold *ad infinitum*. Moreover, some types of personal data require respect for privacy, namely *data volunteered by users*, such as when a social media user creates a profile, and *observed data* that are not actively shared, such as an individual's web search history or location data collected by mobile phones. But identifiable natural persons are not deemed to be owners of data generated in this manner.

Even more controversial is the third type of personal data, namely *machine-generated data* inferred from data analytics, such as credit scores derived from individuals' online payment history.^{iv} Who has the right to control, access, and reuse such data? For example, does the car owner own the data generated by her vehicle, or can the manufacturer lay claim to use and resell the data? There is no single answer from an economic or legal perspective. In reality, some manufacturers such as John Deere and General Motors insist that they own the software embedded in a tractor or a car, and by extension control the data generated by the software.^v In Internet-of things settings, machine generated data often end up under the *de facto* exclusive control of one party because sensors and machines are designed to achieve that outcome.

Thus, the algorithmic use of data creates ownership ambiguities, especially when data is generated in the process of using a service. As a result, we can expect a variety of solutions to the challenges posed by contracting for data.

Allocation of responsibility for training

The next peril in contracting for AI is that it is hard to distribute the reward for insights generated from training a model, along with allocating responsibility when things go wrong. An enterprise software startup with great AI/ML technology

might approach a large corporation, promises to process and analyze large amounts of customer data, so as to fine tune its algorithms for the benefit of this customer.

In legal services, for example, law firms may utilize an ML model which has been pre-trained on publicly available data (this is called an out-of-the-box solution) or trained internally on proprietary data. Lawtech providers typically give clients a choice between the out-of-the-box solution and the train-it-yourself-from-scratch solution. Law firms might ideally wish to buy the out-of-the-box solution and do further training with proprietary data. But this third way is sometimes not offered, due to the so-called “black box problem” that may result in disputes over liabilities.^{vi}

Humans have full visibility over the data (as an input) and the results (the output) that the AI tool consumes or produces. However, what happens in between, especially with deep learning and neural networks, is a black box with little transparency as to how or why the model produced a specific output. Lack of transparency can create complex liability issues. When AI tools and services that function without much human intervention fail, it becomes particularly difficult to determine who bears responsibility. The failure might result from a technical malfunction or from human error in the way in which the model was designed and/or trained. Also, improving an algorithm involves trial and error, making it difficult to attribute the cumulative benefit of adding a new feature (i.e. a measurable characteristic of a phenomenon) to one party or the other, as the existing features may not be orthogonal to each other.

Despite the above challenge, some technology vendors negotiate upfront that while the data would strictly remain the property of the customer, the data learnings from training the model would be owned by the vendor.^{vii} The fear that access to data learnings and data themselves becomes blurred makes corporate clients reluctant to sign such a contract. One resolution is for clients to configure existing solutions internally.

Commercial sensitivity undermining data network effects

A third difficulty in contracting for AI is that commercial sensitivity of data gets in the way of data aggregation. Data-driven learning enables faster and more accurate predictions and recommendations. On the one hand, technology and data analytics vendors would be keen on data pooling to fine tune their algorithm for the benefit of multiple customers. On the other hand, corporate clients with proprietary data remain reluctant to share data across firms. This is particularly the case when the content of data is at the core of a firm’s competitive advantage. Here, the point is that data sharing among businesses is not happening, not because of market failure, but because of corporate strategy.

Strategic reluctance by corporate clients to share data constitutes a barrier to creating data network effects. An AI model exhibits data network effects if the more that it learns from data, the more valuable it becomes to users.^{viii} Value

may be in terms of superior functionalities, such as a more personalized experience for each user. This is distinct from the normal network effect, with the platform utility rising with greater number of users (e.g. of mobile phones). Of course, the two are related as more users → more data → smarter algorithms → better product → more users.

Data network effects – the extent to which data-driven learning exists -- are more important in certain segments of the market than others. For example in fintech, data aggregation is more important in car or health insurance aimed at greater personalization in actuarial science than in investment or retail banking. In lawtech, both a large data volume and a large number of data points lead to greater data-driven learning in contract analytics and legal research than in M&A due diligence or litigation support (e-discovery), both one-off with less benefit from data aggregation.

Data network effects also present an entrepreneurial opportunity.^{ix} Indeed, there exist startups such as Flatiron Health (in oncology-based health data)(acquired by pharma company Roche for \$1.8bn in 2018), Kabbage (in financial services for small businesses)(acquired by American Express in August 2020), and Otonomo (in connected car and mobility data)(with a \$1.4bn valuation at IPO in August 2021). Their central proposition is to extract value from data. To incentivize data-owning customers to buy in, some ventures offer a contributory model: customers have to join the “customer learning network” if they wish to benefit from what the product learned from all other customers. An alternative is a tiered pricing where the customer pays more if they decide to not join the “customer learning network.”

Moving forward

Contracting for AI shares many of the challenges when technology providers are expected to incorporate innovation in products and services. This requires incentivizing and rewarding suppliers for their efforts without being able to specify in advance the exact form of an innovative product or service. In such situations, contracting professionals advocate “relational contracts”,^x in which good communication facilitates developing a collaborative culture of joint problem-solving and “gain and pain sharing”.

This column argued that contracting for AI raises further unresolved issues due to three perils, namely uncertainty over data ownership, assigning benefits and liabilities of training AI models, and commercial sensitivity. Consequently, many companies remain reluctant to share their data. Thus, much data remain locked up and are not available for re-use, undermining opportunities to boost the productivity and innovative activities of firms. Policy recommendations might include the government mandating data pooling in specific areas like health and transportation, as the Finnish government has done.^{xi} Government

procurement of AI-enabled services might also be a lever to enhance good AI contract management.^{xii}

Over time, technological advances may ease the problems. In particular, improvements in explainable AI methods (to ensure that each decision made during the ML process can be traced and explained) would mitigate the “black box” lack of transparency. Also, advances in federated ML would lessen concern for data security and privacy. Nevertheless, these and other technology solutions, e.g. to enhance data portability, are a necessary but not sufficient condition for good contract management. And there remains much scope for aligning incentives to balance vendors’ wish to aggregate data to train their models against their clients’ wish to preserve their commercial edge via proprietary data. But at present, we are at an early stage when viable solutions might take a good deal longer to emerge than we think.

- i World Commerce & Contracting Association (<https://www.worldcc.com/>).
- ii <https://www.microsoft.com/en-us/insidetrack/digitizing-contract-management-using-micro-services-on-azure>
- iii Christine Borgman (2017) *Big Data, Little Data, No Data*, Cambridge, Mass.: MIT Press.
- iv For three types of personal data, see OECD (2015) *Data-Driven Innovation: Big Data for Growth and Well-Being*, Paris: OECD, p.451. Also, Martens, B. et al (2020) *Business-to-Business data sharing: an economic and legal analysis*, JRC Digital Economy Working Paper 2020-05, European Commission, p.42.
- v Kyle Wiens (2015) “We can’t let John Deere destroy the very idea of ownership” *Wired*, April 15.
<https://www.wired.com/2015/04/dmca-ownership-john-deere/>
- vi <https://www.brownejacobson.com/insights/common-ai-related-technology-project-disputes-and-how-to-prevent-them>
- vii <https://mattturck.com/the-power-of-data-network-effects/>
- viii Currier, J. What makes data valuable: The truth about data network effects. NFX.com (Feb. 20, 2020). Also, Gregory, R. W., Henfridsson, O., Kaganer, E., & Kyriakou, H. 2021. The role of artificial intelligence and data network effects for creating user value. *Academy of Management Review*, 46(3): 534-551.
- ix Michael A. Cusumano (2022) “Data platforms and network effects” *Communications of the ACM* 65(10): 22-24.
- x <https://www.worldcc.com/Learn/Workshops/Relational-Contracting>
- xi <https://www.ga4gh.org/news/ga4gh-gdpr-brief-the-finnish-secondary-use-act-2019-may-2020-bonus-brief/>
- xii Coglianese, C., & Lampmann, E. 2021. Contracting for algorithmic accountability. *Administrative Law Review Accord*, 6: 175.