

Final Report on Visiting Scientist (Research) Fellowship of Victoria Lemieux

Fellow's name	Victoria Lemieux
Institutional affiliation	University of British Columbia, School of Information and Institute for Computing, Information and Cognitive Systems
Fellowship period	Jan. 15, 2026 – June 30, 2026
Research project title	Investigating the Affordances of Blockchain as a Juridical System
Report submission date	June 30, 2026

1. Fellowship Objectives

The original aim of my visiting research fellowship was to advance exploration of blockchains as juridical systems, building on chapter 8 of my 2022 book *Searching for trust: blockchain technology in an age of disinformation*¹ in which I argued blockchains and distributed ledgers, like records of all kinds, can be viewed as socio-informational-technical in nature², not just technical artefacts. They function as what is described in archival science as a juridical system - '[a] social group that is organized on the basis of a system of rules and that includes three components: the social group, the organizational principle of the social group, and the system of binding rules recognized by the social group.'³ In the blockchain literature, juridical systems are often characterized as a form of 'digital commons', with Elinor Ostrom's theory of governing the commons being the most prevalent guiding analytic perspective.⁴ Conceptualizing of blockchains and distributed ledgers this way surfaces that they are governed by principles, procedures, and rules created by social actors aimed at arriving on a consensus on the creation of a distributed ledger. These 'epistemic rules of the game', and the version of social 'truth' that results from their operation, are deeply interrelated to epistemic and social trust, and with power and authority. When peer disagreements arise within blockchain and distributed ledger ecosystems, much can be learned from how such epistemic disagreements are resolved (or not) about what version of 'truth' concerning a state of affairs will be accepted in the face of two or more alternate justified true beliefs, and about the basic character and cohesiveness of these ecosystems and the socio-epistemic realities they constitute (Lemieux 2022).

Initially, I intended to focus the above dynamics to extend my work on mis- and disinformation, focusing particularly on questions of resolving peer disagreements; however, given the rise of artificial intelligence and

¹ Lemieux, Victoria L. *Searching for trust: blockchain technology in an age of disinformation*. Cambridge University Press, 2022.

² Lemieux, Victoria L., and Chen Feng. "Conclusion: Theorizing from multidisciplinary perspectives on the design of Blockchain and distributed ledger systems (part 2)." In *Building decentralized trust: Multidisciplinary perspectives on the Design of Blockchains and Distributed Ledgers*, pp. 129-163. Cham: Springer International Publishing, 2021.

³ InterPARES 2 Project. (2021a). Document. In *The InterPARES 2 project glossary*. Retrieved from http://www.interpares.org/ip2/display_file.cfm?doc=ip2_glossary.pdf

⁴ Ostrom E *Governing the commons: the evolution of institutions for collective action*, 2nd edn. Cambridge University Press, Cambridge, 2015.

concerns about AI safety, including concerns that can be described as epistemic in character (e.g., bias, hallucination, and manipulation of beliefs), **I decided to focus instead on the question of whether blockchain-based juridical systems could be used to address AI safety concerns.** Specifically, I saw the governance affordances of blockchains and distributed ledgers as they most often manifest in protocols concerned with the operation of blockchain and distributed ledger ecosystems and Decentralized Autonomous Organizations (i.e., an inward facing governance) as presenting governance capabilities that could be used to govern AI (i.e., an outward facing governance). This perspective constitutes a departure from most of the literature on the use of blockchains to defend against the risks of AI because it focuses on the application of the socio-epistemic character of blockchains and DLTs to defend against AI risks, rather than on their technical affordances (e.g., tamper-proofing images etc by recording hashes of the content on ledger, or ‘tokenization’). I saw this research also as being in keeping with ZIBR’s mission to “explore blockchain as a force reshaping politics, law, culture, and everyday life” (ZIBR website).

2. Research Context and Rationale

Artificial intelligence is becoming embedded in research, public administration, cultural heritage, and everyday decision-making. Yet many of its most immediate risks do not arise from hypothetical superintelligence or ‘existential risks’. Rather, they arise from the large-scale extraction of data, opaque algorithmic processing, weak accountability, concentration of infrastructural power, and the use of automated systems in contexts where privacy, copyright, authenticity, and social trust matter. Archives make these problems especially visible because they hold records created under legal and ethical obligations that often pre-date contemporary machine learning. Thus, my research focused on how blockchain’s governance capabilities could be applied to protect cultural heritage materials from the risks of AI, building upon prior work to develop a blockchain-based solution to protect the privacy of archival documents (Lemieux & Werner, 2023; Lemieux et al, 2024).

3. Research Activities and Methodology

The research focused on developing the Clio-X research programme as a research-through-design intervention combining compute-to-data with blockchain-enabled provenance and integrity controls, auditable transactions, smart-contract-based permissions, and community governance. After an initial review of the literature on AI safety, the analysis mapped these capabilities to five AI safety requirements: robustness, transparency, controllability, alignment, and monitoring. I used this mapping to argue that blockchain is strongest as a governance and trust infrastructure: it can make data access rules more enforceable, document actions, distribute authority, and support accountable oversight. However, blockchain does not solve model interpretability, hallucination, bias, semantic validity, or the problem of aligning autonomous systems with human values. These limitations require human judgment, institutional accountability, privacy-enhancing computation, and participatory governance. Based on this analysis, I rejected both technological solutionism and a narrow view of blockchain as merely a tamper-resistant ledger and proposed a socio-informational-technical model in which archives remain custodians and co-governors of AI infrastructure, rather than becoming passive suppliers of training data. The resulting framework that emerged positions blockchain as one component of a broader, community-centred architecture for responsible AI.

To explore this basic proposition, I undertook a series of interviews with archives, cultural heritage and records professionals at various conferences and online. During the period February to May 2026, I completed 29 separate interviews. In addition, I presented Clio-X as an AI governance infrastructure at 10 online and in person conferences and webinars, listed below, to gain feedback and refine my emerging concepts and assumptions. The iterative combination of design and user research provided an important methodological contribution. My project did not treat users simply as recipients of an already completed technology. Their

concerns were used to identify design requirements relating to usability, transparency, training, governance, meaningful outputs, institutional legitimacy, and community participation. The user-study evidence provides an empirical account of how archivists, researchers, and related professionals understood Clio-X infrastructure. It documented not only functional requirements, but also the social, economic, ethical, and institutional conditions under which such infrastructure might, or might not, be adopted. The list of venues for the presentations is as follows:

- University of Geneva, Geneva
- AUdS Conference, Prague
- DLM Forum, Bern
- DLM Forum, online webinar
- ZIBR Public Talk
- Digital Stewardship Knowledge Exchange, Council of Prairie and Pacific University Libraries, online webinar
- Bibliocon 2026, Berlin
- European Association for Banking History presentation, Milan
- IDAWG, European Database Group, online webinar
- OLOS Foundation, Switzerland, online presentation.

I have documented the outcome of these interviews and workshops in field notes, which I have partially analyzed at this point. The completed analysis will be reported on in two additional planned research papers.

During the period of the fellowship, I have disseminated preliminary findings in a public talk given at University College London on May 7, in the InterPARES Trust AI Summer School and Symposium in held in San Benedetto del Tronto, June 21-26, and in a paper (attached) submitted in June to the journal *Data & Policy* (Cambridge U. Press) written in collaboration Dr. Esther Olembe from the National Archives of Cameroon.

Table 1: Activity Summary Table

Activity	Purpose	Period	Status or Result
Literature review	Refine conceptual framework	January -February 2026	Completed
Stakeholder interviews	Gather empirical evidence	February – May 2026	29 interviews completed; analysis in progress
Presentations and webinars (In-person and online)	Validate preliminary findings & generate new hypotheses	February – June 2026	10 presentations given; findings refined; analysis in progress
Dissemination of Results	Develop research papers and present findings	May-June	1 presentation at UCL in May 2 presentations at the InterPARES Trust AI Summer School & Symposium 1 research paper developed and submitted to <i>Data & Policy</i> (Cambridge U. Press) 2 additional research papers planned

4. Principal Findings and Contributions

The research demonstrated that archives face a fundamental tension between providing access to digital records and protecting the personal, confidential, culturally sensitive, and legally restricted information contained within them in the age of AI. Blockchain and distributed ledger technologies, combined with privacy-enhancing AI, in data space frameworks, such as Clio-X does, offer a promising alternative to the conventional assumption that records must be opened, copied, transferred, or fully reviewed before they can be used. They make it possible to bring approved and governed computational processes to protected data and release only limited or aggregated analytical results. In principle, this enables researchers to learn from restricted collections without gaining direct access to the underlying records.

The user evaluation nevertheless shows that technical feasibility alone does not automatically produce institutional acceptance. Participants generally recognized the value of privacy and sovereignty-preserving access, but adoption was constrained by:

- lack of transparency about how algorithms and the platform operate;
- concerns that the system functions as a technological “black box”;
- uncertainty about the reliability and meaning of analytical outputs;
- the perceived complexity of Web3 and blockchain technologies;
- limited time, funding, skills, and institutional capacity;
- concerns about automation, professional displacement, and the wider effects of AI;
- uncertainty about who would own, operate, finance, and govern the platform; and
- mistrust of commercially driven or promotional claims about AI.

A particularly important finding is that trust is not a one-time decision or an automatic product of technical security. It is a continuing institutional and relational process. Users require visible evidence of transparency, accountability, legitimacy, responsiveness, and ethical conduct throughout the system’s operation. They also prefer to rely upon existing trust relations already established in “designated communities” associated with digital preservation or document sharing (e.g., Europeana).

The research also finds that governance is not a secondary layer added after technical development. Governance is part of the core infrastructure through which privacy, legitimacy, accountability, and user confidence are produced. Participants were more likely to trust a system governed by archives, universities, public-interest bodies, professional communities, or transparent consortia than one controlled solely by a private technology provider.

New concepts, methods, evidence, or datasets

The research develops Clio-X as a new model for governed AI. It conceptualizes archival access not simply as the disclosure of documents, but as infrastructurally controlled (e.g., using the affordances of blockchain technology) execution of computational operations over records that remain under institutional custody. A central concept is compute-to-data. Rather than transferring archival records to researchers or external AI providers, approved algorithms are brought into a protected environment where the records are held. Researchers receive authorized outputs—such as statistics, topic models, visualizations, or aggregate patterns—rather than unrestricted copies of the source records. The research also advances the concept of privacy-preserving distant reading. Distant-reading methods allow researchers to identify broad themes, relationships, trends, sentiments, or communication patterns across large collections without necessarily inspecting individual documents. When combined with controlled computing and disclosure review, this creates a graduated form of access between complete closure and unrestricted disclosure.

Contribution to relevant fields

The research contributes to archival science by reframing the longstanding conflict between access and privacy. Conventional approaches generally assume that archivists must either review and release documents or restrict access to them. Clio-X introduces a third possibility: enabling bounded forms of computational use while the records themselves remain protected. It extends archival thinking beyond privacy protection at rest or during transfer to include privacy during computation. This is increasingly important because AI models and analytical processes may expose, infer, reproduce, or retain sensitive information even when conventional access controls are in place.

More broadly, the research positions archives as active AI-governance institutions. Archives do not merely supply historical data to AI systems; they establish the evidentiary, ethical, and custodial conditions under which records may be used responsibly.

The research also contributes to digital humanities and cultural-heritage computing by proposing an infrastructure through which scholars may analyze otherwise inaccessible collections. This could expand the evidentiary base of historical and cultural research while reducing the need to disclose sensitive source material, although historians interviewed for the study stated that they would still need to review underlying documents to validate analytic insights.

For AI governance, the project shows that responsible AI depends on more than model accuracy or compliance documentation. It also requires rules and infrastructure governing:

- where data remains;
- which algorithms may access it;
- who authorizes computation;
- what outputs may leave the secure environment;
- how actions are logged and audited;
- how affected communities participate; and
- how responsibility is allocated among institutions and technology providers.

The research makes a theoretical contribution to the literature on blockchain governance by reconceptualizing it as a capability that can be applied to AI governance and by extending and challenging conventional innovation-adoption models. Rogers' framework helps explain the relevance of complexity, compatibility, trialability, and observability. However, the findings show that adoption of blockchain-based infrastructure for AI governance cannot be understood primarily as an individual decision based on perceived advantage. Trust is negotiated continuously, while governance, political economy, professional insecurity, institutional legitimacy, and collective accountability are central rather than peripheral factors.

Policy or practice implications

The research highlights that institutions adopting AI-enabled access systems desire to maintain custody and effective control over their records. Contracts and technical arrangements should prevent external providers from retaining records, reusing them for model training, or transferring them to unapproved environments. AI procurement and implementation policies should require privacy and security across the full data and model lifecycle. This includes evaluating whether models, embeddings, intermediate outputs, logs, or visualizations could leak or permit the reconstruction of protected information.

Specific communities could use blockchain-enabled infrastructures to establish juridical systems that articulate AI safety procedures such as:

- approving algorithms and research questions;
- authenticating researchers and assigning permissions;
- assessing the disclosure risk of proposed outputs;
- logging and auditing computational activity;
- responding to incidents and contested results;
- evaluating models for bias, accuracy, and suitability;
- documenting data provenance and transformation; and
- periodically reviewing access and governance decisions.

Overall, research participants confirmed that human oversight remains necessary. Automated tools should initially support rather than replace professional judgment, especially where sensitivity is contextual, culturally specific, reputational, relational, or dependent on information distributed across multiple records. Participants also confirmed that pilot projects should begin with bounded collections, clearly defined research uses, low-risk analytical outputs, and measurable evaluation criteria. Users need opportunities to test the system and observe its safeguards before being expected to trust or adopt it. They also conveyed that training should address not only how to operate the technology but also how to interpret computational results, recognize limitations, assess disclosure risk, and understand the governance arrangements behind the platform. Participants' reactions confirmed that explanations of how Clio-X serves to govern AI should avoid unnecessary Web3 terminology and focus on recognizable archival tasks and benefits. Governance arrangements should include archivists, researchers, technical experts, privacy and legal specialists, records creators, and communities represented in the collections. Feedback received during interviews and workshops indicated that public-interest partnerships, academic consortia, or other community-based structures are viewed as providing greater legitimacy than purely commercial control. Study participants also indicated that policy should also recognize institutional capacity constraints. Without sustainable funding, staffing, infrastructure, and technical support, Clio-X could increase rather than reduce burdens on archival institutions.

Limitations of the research

The Clio-X user evaluation involved a relatively small number of workshops and participants. Although the participants represented different locations, professional roles, and perspectives, the results cannot be assumed to represent the archival or digital-humanities communities as a whole. Researchers and other prospective data consumers were underrepresented. Consequently, the study provides stronger evidence about the concerns of archival and professional stakeholders than about researchers' willingness to use constrained analytical access or the usefulness of the outputs for actual scholarship.

The workshop time also was limited, and some participants could not respond fully to every question. Complex concepts such as data spaces, blockchain, trusted execution environments, privacy-preserving machine learning, and decentralized governance may also have been difficult to assess without extensive practical experience.

The technical analysis establishes potential and conceptual feasibility more strongly than real-world effectiveness. Further empirical evidence is required concerning computational performance, cost, interoperability, scalability, disclosure risks, usability, and compliance in operational archival settings.

The effectiveness of the system also depends on the quality of source data, metadata, digitization, optical character recognition, and archival description. Privacy-preserving computation cannot correct incomplete, biased, poorly described, or inaccurately digitized records automatically.

Finally, blockchain or decentralized governance does not by itself guarantee legitimate, inclusive, or equitable decision-making. Token-based participation, smart-contract voting, and decentralized autonomous organizations may reproduce inequalities if authority, expertise, resources, or voting power are distributed unevenly.

Questions requiring further investigation

My future research will investigate the following specific questions as a follow-on from this research:

1. How does Clio-X as an AI governance infrastructure perform when deployed with real archival collections, institutional workflows, researchers, and access requests?
2. Which combinations of privacy-enhancing technologies provide the most appropriate balance among privacy, analytical utility, cost, speed, and interpretability for different archival uses?
3. How should institutions assess the disclosure risks of aggregate outputs, visualizations, model parameters, embeddings, and repeated queries?
4. Which forms of analysis can be safely supported without prior document-level sensitivity review, and which continue to require direct human examination?
5. How useful are privacy-preserving outputs for genuine historical, social-scientific, genealogical, legal, Indigenous, or community-based research?
6. How do researchers adapt their questions and methods when they cannot inspect or download the underlying records?
7. What technical and evidentiary information must be made visible for users to understand and trust the system without overwhelming them with complexity?
8. What governance model is most appropriate: direct archival control, a public consortium, an academic partnership, community governance, a decentralized autonomous organization, or a hybrid arrangement?
9. How should people and communities represented in archival records participate in decisions about computational access and acceptable uses?
10. How can accountability be allocated among archives, software developers, AI providers, infrastructure operators, researchers, and governance bodies when harm or unauthorized disclosure occurs?
11. What legal and regulatory frameworks apply when algorithms cross institutional or national boundaries while the data remains in its original jurisdiction?
12. How can the architecture preserve archival authenticity, provenance, chain of custody, and interpretive context throughout computational processing?
13. What skills, staffing models, procurement practices, and sustainable funding arrangements are required for smaller or under-resourced institutions to participate?
14. Does community-based or decentralized governance measurably increase trust, and under what conditions might it instead introduce new complexity or inequality?
15. How should the success of privacy-preserving archival access be measured—including increased access, reduced review costs, research value, privacy protection, user confidence, institutional legitimacy, and benefits to represented communities?

5. Engagement and Collaboration with ZIBR

Address the following points as relevant:

During my time at ZIBR, I also was able to engage and collaborate with ZIBR chairs, researchers and staff in numerous ways. For example, I was able to participate as a panelist on Tech-powered Governance: Law, Trust, and Institutions during the launch of ZIBR in January. I also organized two ZIBR public talks, one asking whether Crypto could Rebuild the Scientific Research Process and Protect Intellectual Assets with Dr. Luis de la Torre from UNED in Madrid, and another on Retail Central Bank Digital Currency with Dr. Geoffrey Goodell from University College London. In addition, I participated in a ZIBR event on “Holding Volatility Workshop: Perspectives on Exocapitalism,” organized by fellow Visiting Research Scientist, Dr. Andrea Leiter. I also contributed to two publications in my capacity as a ZIBR Visiting Research Scientist: The UNILU Magazine, to which I contributed a piece with Dr. Andrea Leiter entitled “Wandel Von Wissen und Regeln” (attached) and a contribution to the 2026 CV VC Top 50 Report focusing on the connection between archival science and blockchains (attached).

In addition, to the above, I engaged with UNILU faculty Mira Buri and Sophie Mutzl, incoming Visiting Research Fellow Annika Aebli, external organizations such as the Office of the Canadian Trade Commissioner, British Blockchain Association as an invited guest of their annual scientific conference held at the UK Houses of Parliament, the Swiss FinTech Association, which held an event at CV Labs on predication markets that has inspired me to explore the epistemological dimensions of these markets as a future research topic. I further engaged with Global SME Forums and was able to facilitate the participation of ZIBR Research Chair, Dr. Adam Hayes as a panelist on the future of technology at their event in Davos on June 4, an event in which I also participated. Additionally, I participated as an expert in the plenary meeting of the International Standards Organizations’ Technical Committee on Blockchain and Distributed Ledger Technology, for which I have served as Project Leader for ISO 22739 on Blockchain and DLT Vocabulary since 2017 and as Head of Delegation for the Canadian member body. These engagements contributed to developing ZIBR’s wider research and industry network.

I also developed new research relationships within ZIBR and connected ZIBR Chairs with my own wider research network to lead the development of a Letter of Intent for the New Frontiers in Research Call for Proposals on Disruptive Technology (attached). Our proposed research consortium comprises multi- and interdisciplinary researchers from Canada, Finland, Ghana, the Netherlands, Sweden, and Switzerland (ZIBR and HSLU) and will focus on how data spaces can make sensitive and fragmented data usable for SDG action without undermining privacy, sovereignty, rights, or trust in health, energy, biodiversity and culture. Over the summer we will peer review other Letters of Intent for this CFP and expect to hear whether we will be invited to submit a full proposal by the end of August.

Finally, I contributed to the future development of ZIBR by providing feedback on a new policy on Visiting Research Fellows, reviewing applicants for 2026 Visiting Research Fellows, and by developing in collaboration with Dr. Andrea Leiter a new policy on Former Visiting Research Fellows and Visiting Fellows (attached).

6. Appreciation for the Fellowship Experience

I would like to end this report by saying how thoroughly I appreciated my fellowship time at ZIBR. During my stay, I was welcomed into the ZIBR community and provided with excellent administrative support, facilities and opportunities to engage and collaborate. I think one aspect of my time at ZIBR that I enjoyed the most was the informal discussions held over lunch with ZIBR Chairs, researchers and staff about our respective research, societal developments and broader topics of general interest. I hope that this tradition continues.

My sincere thanks to the ZIBR leadership for affording me the opportunity to spend time with you and to Martina Kunz, ZIBR Administrator for excellent, efficient and welcoming administrative support. I will briefly

allude here to some of the personal challenges I faced during my time at ZIBR. I credit the warm and collegial atmosphere at ZIBR as helping to give me the fortitude to face and overcome these challenges. I leave ZIBR feeling that my fellowship allowed me the time and inspiring environment to grow in my scholarship and share my experience as a senior academic researcher, but also that at six months in length, despite accomplishing a good amount, my fellowship was too short. I would have loved to stay for an entire year, but alas!

Appendices

Address the following points as relevant:

- Paper submitted to *Data & Policy* (Cambridge U Press) – “Making AI Governable: Clio-X, Public Records, and Community-Based Digital Statecraft in Cameroon”
- Article for UNILU Magazine - “Wandel Von Wissen und Regeln”
- Article for CV VC Top 50 Report
- NFRF LOI – “Resilient Data Infrastructure to Advance the Trustworthy Application of AI and the SDGs”
- ZIBR Policy on Former Visiting Fellows and Researchers
- Photos of Events