Governing the Whole Stack: Auditable Data Science Queries for Frugal and Sovereign Environmental AI

Genoveva Vargas-Solar[1]

[1]CNRS, Univ Lyon, INSA Lyon, UCBL, LIRIS, UMR5205, F-69221, France

Abstract

Environmental AI increasingly relies on end-to-end data science pipelines that run across distributed infrastructures such as edge, fog, and cloud, and involve multiple stakeholders, including public agencies, communities, researchers, and service providers. However, these pipelines are still mainly evaluated through technical criteria such as accuracy and latency, while important concerns such as environmental cost, fairness, and sovereignty remain difficult to specify, monitor, and enforce. This paper introduces *auditable Data Science Queries* (DSQs), a whole-stack abstraction that treats data science pipelines as executable queries with explicit contracts. In addition to performance and reliability, these contracts include constraints on energy and carbon cost, fairness, and data/model/compute sovereignty. We propose a layered architecture that translates governance decisions into machine-actionable policies for resource placement, scheduling, and execution, while provenance and observability mechanisms generate evidence for auditing and contestability. We also define a fairness index and preference model to guide runtime resource allocation under multiple objectives. A river-basin use case for water quality monitoring and equitable allocation illustrates how governance requirements can shape execution decisions and produce accountable evidence.

Keywords

Auditable Data Science Queries, Whole-stack governance, Frugal/sustainable Environmental AI, Data/model/-compute sovereignty, Provenance & accountability

1. Introduction

The rapid growth of AI-based data analytics has increased the scale of scientific computing, but also its environmental, economic, and social costs. Training and deploying models now require large amounts of energy, computing infrastructure, and money. At the same time, many datasets and models remain incomplete, biased, and dominated by digital content produced in English and in the Global North, which can marginalise other knowledge systems and reinforce existing inequalities. These dynamics raise broader concerns about technological dependency, concentration of power, and new forms of digital colonialism.

Much of the infrastructure behind data analytics remains invisible to users: where data is stored, where computation takes place, who controls the platforms, and under which governance rules. This raises important questions about data origin, consent, localisation, and accountability. When data concerns specific territories, communities, or individuals, it is necessary to ask not only whether they agreed to its use, but also whether they are aware of how it is processed, where the infrastructure is located, and what environmental cost is incurred.

Our work studies resource allocation in *Data Science Query* (DSQ) environments,[1] where scientific data-driven processes are executed as end-to-end computational workflows. We focus on how to allocate resources while satisfying both classical service-level objectives (SLOs) and broader fairness requirements, including server location, data provenance, sovereignty, energy use, carbon footprint, and economic cost. Our objective is not only to optimise execution, but also to document and justify how resources are selected and used, so that these decisions become accountable and auditable. In

Published in the Proceedings of the Workshops of the EDBT/ICDT 2026 Joint Conference (March 24-27, 2026), Tampere, Finland.
genoveva.vargas-solar@cnrs.fr (G. Vargas-Solar)

[1]A Data Science Query (DSQ) is an end-to-end executable pipeline for data access, curation, representation building, training/inference, evaluation, and reporting, packaged with explicit constraints and guarantees.

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-04-09

this sense, the paper contributes to a system-level framework for responsible data science, in which socio-political values are translated into computable and contestable mechanisms.

Contributions. This paper proposes an environment and a methodology that combine technical mechanisms, such as provenance, governed lakehouse/data-lake architectures, multi-objective resource allocation, and audit protocols, with a social-sciences perspective grounded in STS, feminist epistemologies, decolonial approaches, and environmental justice. In our context, a decolonial perspective means supporting sovereignty over data, models, and computing resources, including the possibility of adapting, reusing, and governing technologies in ways that respect local needs and constraints.

The paper introduces *Auditable Data Science Queries* (DSQs) for environmental AI, proposes a whole-stack governance architecture to translate participatory decisions into executable policies, defines a fairness-aware dispatching and negotiation mechanism for runtime resource allocation, and demonstrates the approach through a river-basin use case for water quality monitoring and equitable allocation.

The remainder of the paper is organised as follows. Section 2 reviews the related literature. Section 3 presents the vision of inclusive, fair, and sustainable data science and introduces the whole-stack architecture. Section 4 describes the methodology for designing equitable and locally accountable data-driven systems. Section 5 details our approach to fair DSQ execution. Section 6 presents the river-basin use case. Section 7 concludes the paper and outlines future work.

2. Related Work

This work brings together several research strands that motivate a *whole-stack* view of environmental AI, in which knowledge production, governance, and execution constraints are treated as interdependent system properties.

First, critical algorithm studies, feminist data scholarship, STS, and decolonial AI show that data and models are not neutral: they reflect situated assumptions, power relations, and uneven infrastructures [1, 2, 3, 4, 5, 6, 7]. This literature motivates treating sovereignty, accountability, and epistemic plurality as design requirements rather than as external ethical concerns. Related work on digital sovereignty and commons further shows that control over data, models, and compute depends on technical, legal, and institutional arrangements [8, 9, 10].

Second, environmental justice, sustainable machine learning, and responsible AI provide concepts and tools for extending system evaluation beyond performance. Environmental justice frames fairness in distributional, procedural, and recognition terms [11, 12], while sustainable ML makes energy and carbon visible as optimisation objectives [13, 14]. Responsible AI contributes documentation mechanisms such as datasheets and model cards to make assumptions, provenance, and limitations explicit and contestable [15, 16]. At the infrastructure level, current cloud and distributed systems increasingly expose sustainability metrics and support execution across cloud, edge, and federated settings, but access and control remain uneven, especially in resource-constrained or politically dependent contexts [17, 18, 19].

Third, systems research on fairness, resource allocation, provenance, and auditable execution provides the technical basis for our proposal. Prior work studies fairness in datasets and analytics pipelines [20, 21], fair allocation of compute resources [22, 23], and provenance and reproducibility in scientific workflows [24, 25]. Additional work on black-box trust, certification, and accountable data systems highlights the need for runtime signals that can support auditable decisions under partial visibility [26, 27, 28]. Our contribution builds on these foundations, but differs in combining them into a single framework where fairness, frugality, sovereignty, and accountability are encoded as explicit execution constraints for data science queries.

In summary, the literature shows the need to move beyond accuracy- and latency-centred views of data science systems. Our work responds by focusing on three connected challenges: making resource use accountable, enabling fair and transparent allocation across heterogeneous infrastructures, and

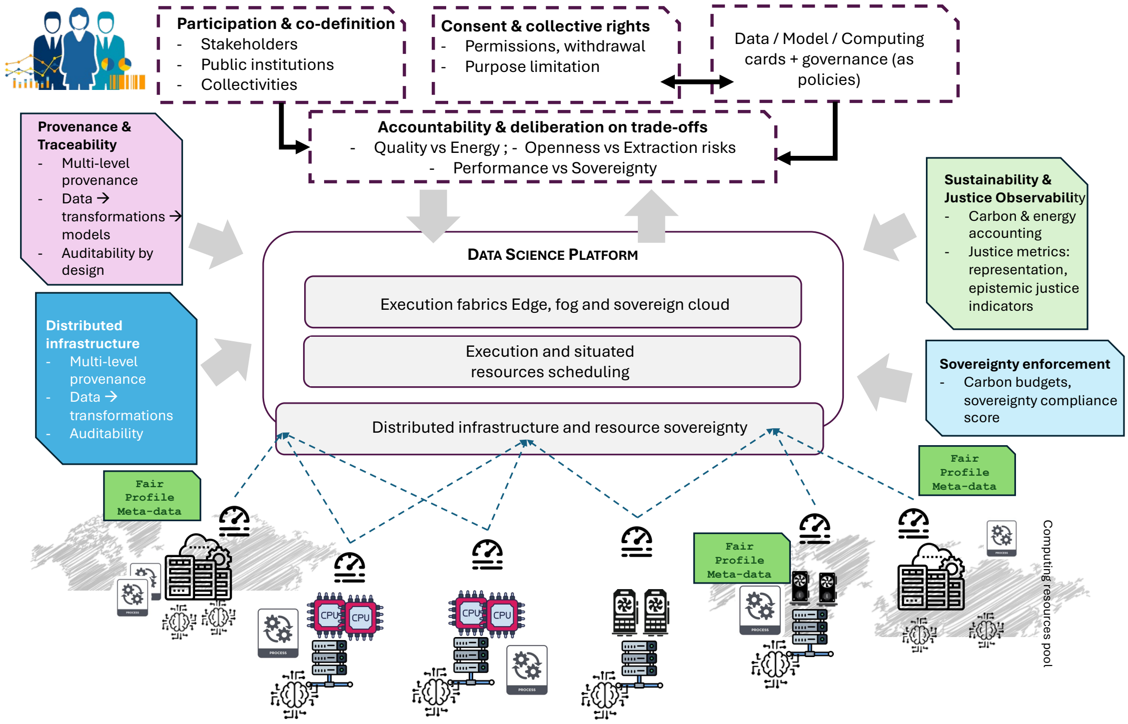

Figure 1: Whole-stack DSQ execution environment where governance and qualitative constraints steer distributed execution, observability, and enforcement.

designing data science queries whose execution remains aligned with the values, rights, and constraints of the communities involved.

3. Towards inclusive, fair and sustainable data science

Making socio-ecosystem data science *responsible by design* requires moving beyond evaluation based only on accuracy and latency. In this work, a data science query is assessed through a combination of performance, environmental and economic cost, fairness, and sovereignty requirements. We therefore define a *Data Science Query* (DSQ) as an end-to-end executable pipeline—including data access, curation, training or inference, evaluation, and reporting—packaged with explicit constraints and guarantees. In addition to classical service-level objectives (SLOs), DSQs include first-class constraints on fairness, energy/carbon footprint, and sovereignty, such as localisation, purpose limitation, and sharing boundaries.

3.1. A layered whole-stack architecture

Figure 1 summarises the proposed whole-stack execution environment. Its main idea is that non-technical requirements, such as consent, rights, and governance choices, are translated into machine-actionable constraints that guide technical execution across edge, fog, and sovereign cloud infrastructures.

At the upper layers, stakeholders co-define the purpose of the DSQ, acceptable uses of data, and the main limits on sharing, retention, and reuse. These decisions are translated into structured artefacts, such as data/model/compute cards and governance policies, that can be validated and enforced during execution. A deliberation layer makes trade-offs explicit, for example between quality, energy, and sovereignty, by turning collective decisions into optimisation parameters.

At the technical core, the platform executes DSQs across heterogeneous infrastructures and uses policy-driven scheduling to allocate compute, storage, and bandwidth under multiple constraints. These include performance goals, fairness preferences, carbon budgets, and sovereignty requirements. A distributed resource substrate manages placement, certification, identity, and secure coordination across sites, ensuring that data and models move only under permitted conditions.

Provenance, observability, and enforcement complete the architecture. Provenance records how data, transformations, models, and results are connected, so that each DSQ run produces not only outputs but also an auditable record of its conditions of production. Observability collects runtime indicators such as energy use, carbon cost, and fairness signals, which are checked against DSQ constraints. Enforcement mechanisms then apply hard constraints, such as localisation and access control, together with softer priorities, such as carbon budgets or sovereignty scores, so that governance decisions effectively shape execution.

Finally, each participating site exposes a *fair profile* describing properties such as location, operator provenance, certifications, hardware profile, and energy characteristics [27]. These profiles connect governance to scheduling by allowing the system to determine which sites are eligible and which ones best satisfy the declared priorities of a DSQ.

3.2. Use case: a DSQ for community-governed water quality

To illustrate the architecture, consider a DSQ that produces weekly water-quality risk maps for a river basin managed by local communities, a public water agency, and researchers. The DSQ combines sensor data, satellite data, and community observations, trains a forecasting model, and produces predictions together with an auditable report.

First, communities and the public agency define the purpose of the DSQ (for example, early warnings for safe water access) and set basic rules: raw community data must stay local, identifiable data cannot be transferred, and results must include uncertainty and limitations. These decisions are then translated into cards and policies that specify purpose, retention, provenance, sharing conditions, and a carbon budget.

Next, these governance choices become execution parameters. For example, stakeholders may decide that locality is mandatory, while carbon reduction is important but secondary during drought periods. The scheduler then uses these priorities to decide where tasks should run. Local feature extraction can run on trusted edge nodes, aggregation can run in a fog layer managed by the public agency, and heavier training can run in a sovereign cloud only when policy allows it. Only approved intermediate results are allowed to cross boundaries.

During execution, the system records provenance, including which data were used, which transformations were applied, where tasks were executed, and which policies shaped the run. It also monitors energy use, carbon cost, and representation indicators. If a carbon budget is exceeded, the system must either renegotiate the execution priorities or switch to a lower-energy plan, while still respecting locality constraints.

This example shows that fairness, sustainability, and sovereignty are treated as executable constraints, not as external principles. In the proposed architecture, participatory decisions become policies, policies guide optimisation, and execution produces evidence that supports auditing and accountable revision.

4. Designing equitable, sustainable, and locally accountable data and algorithm-driven systems

Designing equitable, sustainable, and locally accountable data systems requires governance and execution to be designed together. In this work, we organise the methodology around three connected pillars: *(i)* territorial and epistemic data governance, *(ii)* responsible and frugal model training, and *(iii)* community-in-the-loop resource dispatching. The process is iterative: stakeholders define constraints and guarantees, these are translated into machine-actionable policies, DS queries are executed under multi-objective optimisation, and compliance is checked through observability and provenance.

Backbone use case: equitable water distribution. We illustrate the methodology with a running example in urban-rural water distribution. The goal is to reduce shortages in underserved zones, detect leaks, improve energy efficiency, and maintain local accountability. The DS query combines infrastructure telemetry, consumption data, maintenance logs, and community reports, and produces a weekly decision package including maintenance priorities, inequity risk maps, pressure recommendations, and an auditable report.

Step 1: Territorial and epistemic data governance. The first step defines a community-centred governance envelope specifying what data may be processed, where, by whom, and for which purpose. Datasets are associated with metadata such as consent, locality, retention limits, and sensitive attributes, and these metadata constrain both access and execution. In the water example, this means that detailed community reports may remain local, only aggregated indicators may be shared, and retention may be limited unless renewed collectively, while less sensitive telemetry can be shared more broadly.

Step 2: Frugal and context-aware training. The second step constrains model development through explicit limits on energy, carbon, cost, and fairness. The objective is to avoid unnecessary centralisation and excessive computation while preserving useful performance. In the water example, this favours balanced sampling, lightweight local models, and incremental retraining, with larger models used only when justified and under explicit budgets. Acceptable trade-offs, such as slightly lower accuracy for much lower energy use, are defined in advance with stakeholders.

Step 3: Negotiated resource dispatching. The third step turns governance and training constraints into runtime decisions. DS workflows are executed through a negotiated dispatching process that balances latency, energy, locality, and fairness. In the water example, sensitive feature extraction can run in a local enclave, edge inference can support pumping decisions, and only permitted aggregates are sent to a sovereign municipal cloud. When objectives conflict, the scheduler follows the priority order defined by the community, for example treating locality as non-negotiable while allowing some flexibility on latency or carbon.

Continuous accountability. The methodology is closed-loop: execution produces both outputs and an auditable record of how decisions were made. Provenance and logging document which policies were applied, which data were excluded, which trade-offs were invoked, and what energy and carbon were consumed. In the water use case, this creates a weekly accountability bundle that allows communities and public agencies to verify whether locality was respected, underserved zones were represented, and energy savings were achieved without increasing inequity. This feedback can then be used to revise policies, retraining rules, and scheduling priorities over time.

5. Integrating Fairness into DS Query Execution

Integrating fairness into the execution of data science (DS) queries means treating requirements such as sovereignty, provenance, and environmental impact as explicit optimisation and accountability goals. This requires three elements: a clear definition of what acceptable execution means, measurable indicators that can be monitored during runtime, and execution contracts that make guarantees explicit, auditable, and enforceable.

A DS query involves data, models, and computing resources, all of which may have fairness-related properties, such as location, provenance, energy cost, and governance conditions. For this reason, DSQ execution must consider not only performance and cost, but also where data and models are processed, which resources are used, and whether the execution respects declared constraints. As a result, DSQ outputs should include both analytical results and a structured record of the conditions under which they were produced, including placement decisions, provenance, and evidence of compliance.

To support this, we define a *Fairness Index* (FI), which combines several execution-relevant metrics into a single transparent optimisation objective. The purpose of FI is not to reduce fairness to one universal number, but to make trade-offs explicit and configurable according to the priorities of a community or application. The FI includes metrics related to infrastructure location and provenance, data sovereignty, model performance, training effort, and resource costs.

$$\begin{aligned} \mathrm{FI} = {} & \alpha_1 L_s + \alpha_2 P_s + \alpha_3 DS_p \\ & + \alpha_4 MP + \alpha_5 T_t + \alpha_6 C_g \\ & + \alpha_7 C_c + \beta_1 TC_{CO_2} + \beta_2 EC \end{aligned} \tag{1}$$

Here, L_s denotes server location, P_s server provenance, DS_p data sovereignty and provenance constraints, MP model performance, T_t training time, C_{gpu} computing resources used, C_{cal} calibration cycles, TC_{CO_2} carbon cost, and EC economic cost. The weights encode application-specific priorities and make the optimisation criteria explicit.

Each DSQ also includes preferences, such as admissible locations, acceptable operators, or certification requirements, together with weights that reflect their relative importance. These preferences guide resource selection and placement. Because participating servers may operate as black boxes, monitoring combines *reported* metadata (e.g., certifications, energy profiles, hardware characteristics) with *inferred* signals derived from observable behaviour (e.g., reliability, update cadence, convergence patterns). This makes it possible to assess fairness-relevant properties without direct access to local data.

Finally, resource allocation is treated as a negotiation-aware optimisation problem. The runtime must jointly satisfy classical service-level objectives and fairness constraints expressed by FI. When these objectives conflict, the system computes a best-effort allocation under the declared priorities and records the resulting compromises. In this way, fairness, trust, and accountability become runtime properties of DSQ execution, rather than external concerns.

6. Extended Use Case: River-Basin Water Quality and Equitable Allocation as an Auditable DSQ

We illustrate the approach with a river-basin use case for water quality monitoring and fair water allocation. The goal is to show how the proposed methodology can connect participatory governance, DSQ contracts, frugal modelling, sovereignty-aware execution, negotiated dispatching, and auditability in a single end-to-end scenario.

A river-basin authority coordinates the system across municipalities, rural and Indigenous communities, agricultural cooperatives, and industrial operators. The DSQ combines heterogeneous data sources to produce four outputs: a contamination-risk map, an allocation recommendation, an intervention plan, and an accountability bundle describing the conditions under which these results were produced.

The workflow follows the methodology introduced earlier. First, stakeholders define a governance envelope that specifies locality, purpose limitation, retention, sharing boundaries, and contestability. For example, sensitive community data may remain local and non-exportable, while only certified aggregates may be shared. Second, model development is constrained by fairness and frugality requirements, favouring lightweight local models, sovereign regional processing, and rare heavy retraining under explicit carbon and cost budgets. Third, these constraints are compiled into hard and soft execution rules that guide dispatching across edge, fog, and sovereign cloud infrastructures. When constraints cannot all be satisfied, the system triggers a negotiation process that records any authorised relaxation through an explicit compromise certificate.

Each run produces both domain outputs and an evidence bundle containing policy versions, placement decisions, provenance traces, fairness indicators, carbon and cost reports, and, when needed, compromise certificates. This makes the process inspectable and contestable: communities, agencies, and auditors can verify whether locality was respected, which trade-offs were made, and how these choices affected the final recommendations.

Overall, the use case shows that governance is not external to execution. It determines which resources are admissible, how data and models may move, and how trade-offs are handled. In this sense, the use case demonstrates the practical value of treating fairness, frugality, sovereignty, and accountability as executable properties of DSQ execution.

7. Conclusions and Future Work

This paper advanced the claim that environmental AI demands *whole-stack governance*: fairness, frugality, and sovereignty must be specified as executable constraints and not treated as external checklists. We proposed *auditable Data Science Queries* (DSQs) as a unifying abstraction coupling end-to-end pipelines with explicit contracts on performance and reliability, extended with first-class constraints on energy/carbon, fairness, and data/model/compute sovereignty. We described a layered architecture in which participatory governance is translated into machine-actionable cards and policies that steer distributed scheduling and placement across edge–fog–sovereign-cloud fabrics, while provenance and observability generate audit-ready evidence of how results were produced and which trade-offs were invoked. A multi-metric fairness index and preference model illustrate how community priorities can be compiled into dispatching objectives. The equitable water distribution backbone demonstrated step-by-step how local rights and situated requirements become concrete execution decisions with inspectable accountability bundles.

Future work. We identify five directions that directly extend "governing the whole stack." First, we will formalise DSQ contracts as a typed specification language with verifiable compilation into execution plans and monitors, enabling static checks (policy satisfiability, leakage risk) and runtime conformance checks. Second, we will develop negotiation-aware scheduling algorithms that output minimal and explainable compromise certificates, and we will study their governance properties (who authorises what, under which accountability mechanisms). Third, we will strengthen sovereignty enforcement through continuous certification of data, algorithms, and infrastructures under change [27], linking certification states to scheduling admissibility. Fourth, we will expand observability from carbon accounting to end-to-end "justice observability" (coverage, representation, contestation traces) and evaluate whether these signals effectively support participatory revision. Fifth, we will conduct longitudinal participatory deployments on socio-ecosystem case studies (water distribution, heatwave response, biodiversity monitoring) to evaluate not only predictive performance but governance outcomes: contestability, time-to-revision, dependency reduction, and measurable improvements in energy and equity without increasing infrastructural lock-in. These directions position DSQs as a systems research agenda for auditable, frugal, and sovereign environmental AI.

Acknowledgements This work has been partially funded by the projects FRIENDLY (http://www.vargas-solar.com/friendly) and INFILTRATE (http://www.vargas-solar.com/infiltrate) of the AT program of the LIRIS Lab and the Feminist AI Network https://iafeminista.lat.

Declaration on Generative AI We hereby state that we have used LLM help for producing latex tables adjusted to the size of the two column paper. We used Grammarly premium for verifying English.

References

[1] R. Benjamin, Race After Technology: Abolitionist Tools for the New Jim Code, Polity Press, 2019.
[2] S. U. Noble, Algorithms of Oppression: How Search Engines Reinforce Racism, NYU Press, 2018.
[3] C. D'Ignazio, L. F. Klein, Data Feminism, MIT Press, 2020.
[4] D. Haraway, Situated knowledges: The science question in feminism and the privilege of partial perspective, Feminist Studies 14 (1988) 575–599.
[5] M. Fricker, Epistemic Injustice: Power and the Ethics of Knowing, Oxford University Press, 2007.

[6] N. Couldry, U. A. Mejias, Data colonialism: Rethinking big data's relation to the contemporary subject, Television & New Media 20 (2019) 336–349. doi:10.1177/1527476418796632.

[7] S. Mohamed, M.-T. Png, W. Isaac, Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence, Philosophy & Technology 33 (2020) 659–684.

[8] J. Pohle, T. Thiel, Digital sovereignty, Internet Policy Review 9 (2020). doi:10.14763/2020.4.1532.

[9] L. Floridi, The fight for digital sovereignty: What it is, and why it matters, especially for the EU, Philosophy & Technology 33 (2020) 369–378. doi:10.1007/s13347-020-00423-6.

[10] S. Couture, S. Toupin, What does the notion of "sovereignty" mean when referring to the digital?, New Media & Society 21 (2019) 2305–2322. doi:10.1177/1461444819865984.

[11] R. D. Bullard, Dumping in Dixie: Race, Class, and Environmental Quality, Westview Press, 1990.

[12] D. Schlosberg, Defining Environmental Justice: Theories, Movements, and Nature, Oxford University Press, 2007.

[13] P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, J. Pineau, Towards the systematic reporting of the energy and carbon footprints of machine learning, Journal of Machine Learning Research 21 (2020) 1–43. URL: https://www.jmlr.org/papers/v21/20-312.html.

[14] R. Schwartz, J. Dodge, N. A. Smith, O. Etzioni, Green AI, Communications of the ACM 63 (2020) 54–63. doi:10.1145/3381831.

[15] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumé III, K. Crawford, Datasheets for datasets, arXiv, 2018. doi:10.48550/arXiv.1803.09010. arXiv:1803.09010.

[16] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, T. Gebru, Model cards for model reporting, in: Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19), Association for Computing Machinery, 2019, pp. 220–229. doi:10.1145/3287560.3287596.

[17] E. Commission, European open science cloud strategic implementation plan (2020). Https://ec.europa.eu/research/openscience/pdf/eosc-strategic-implementation-roadmap.pdf.

[18] P. Kairouz, B. McMahan, et al., Advances and open problems in federated learning, Foundations and Trends® in Machine Learning 14 (2021) 1–210. doi:10.1561/2200000083.

[19] W. Shi, J. Cao, Q. Zhang, Y. Li, L. Xu, Edge computing: Vision and challenges, IEEE Internet of Things Journal 3 (2016) 637–646. doi:10.1109/JIOT.2016.2579198.

[20] M. Drosou, H. V. Jagadish, E. Pitoura, J. Stoyanovich, Diversity in big data: A review, Big Data 5 (2017) 73–84.

[21] B. Catania, G. Guerrini, C. Accinelli, Fairness & friends in the data science era, AI & SOCIETY 38 (2023) 721–731.

[22] S. Jiang, J. Wu, Multi-resource allocation in cloud data centers: A trade-off on fairness and efficiency, Concurrency and Computation: Practice and Experience 33 (2021) e6061.

[23] H. Hamzeh, Fairness for Resource Allocation in Cloud Computing, Ph.D. thesis, Bournemouth University, 2021.

[24] Y. L. Simmhan, B. Plale, D. Gannon, A survey of data provenance in e-science, SIGMOD Record 34 (2005) 31–36. doi:10.1145/1084805.1084812.

[25] PROV-DM: The PROV Data Model (W3C Recommendation), Technical Report, World Wide Web Consortium (W3C), 2013.

[26] S. Romdhani, G. Vargas-Solar, N. Bennani, C. Ghedira-Guegan, Qos-based trust evaluation for data services as a black box, in: IEEE International Conference on Web Services (ICWS), IEEE, 2021, pp. 476–481.

[27] M. Anisetti, C. A. Ardagna, N. Bena, Continuous certification of non-functional properties across system changes, in: Proceedings of the International Conference on Service-Oriented Computing (ICSOC 2023), Rome, Italy, 2023. November–December 2023.

[28] C. A. Ardagna, N. Bena, N. Bennani, C. Ghedira-Guegan, N. Grecchi, G. Vargas-Solar, Revisiting trust management in the data economy: A roadmap, IEEE Internet Computing 28 (2024).