Statistical Science for a Better World.

JUNE 2026

vISIon

The Magazine of
The International Statistical Institute

OFFICIAL STATS

Race, ethnicity and data:
what governments choose
to measure, and why it matters.

DATA FOR ALL

Building statistical literacy
for everyone, everywhere,
at every stage of life.

REDUCED INEQUALITIES

Data for a fairer world:
who counts, and
who is counted.

IN THIS ISSUE

Ethnicity & official data • Statistical Literacy • AI & slavery

Policies & Information

Conflicts of interest. Unless otherwise stated, authors declare no competing interests.

Submissions & enquiries. Pitches, letters and article proposals: manuele.leonelli@ie.edu.

Disclaimer. *vISIon: The ISI Magazine* is not a peer-reviewed journal. Views expressed are those of the authors and do not necessarily reflect the International Statistical Institute, its associations, editors, partners or supporters. While care is taken, ISI accepts no responsibility for errors or omissions; readers should assess the appropriateness of any methods described.

About *vISIon*

vISIon: The ISI Magazine is an initiative of the International Statistical Institute for accessible, rigorous writing on pressing societal issues through a data lens. Built on the society's values of professionalism, truthfulness, integrity and respect, the magazine bridges academia, national statistical offices and industry. Each issue centers on one UN Sustainable Development Goal and gathers perspectives from across ISI's seven associations. Our aim is clear, ethical communication that helps society make informed, evidence–based choices.

Copyright

Table of Contents

June 2026 • Vol. 1 No. 2

International Statistical Institute

The Bigger Picture

A THEMATIC OVERVIEW

This issue of *vISIon* explores how statistics helps make hidden realities visible. Across topics ranging from modern slavery, race and ethnicity, and community-engaged research to climate vulnerability, official statistics, education, and computational tools for human rights, the contributions show how data can support accountability, understanding, and better decision-making. Together, they illustrate the many ways statistical thinking contributes to a fairer and more informed society.

Making the invisible visible

VISIBILITY, ACCOUNTABILITY, AND EVIDENCE.

Adriana Bora's feature article on AI against modern slavery explores how explainable artificial intelligence and large-scale document review can support accountability across global supply chains. Jonathan Auerbach's contribution on race, ethnicity, and official statistics examines how classifications shape visibility, recognition, and public policy. Suryo Adi Rakhmawan and I Nyoman Setiawan extend this discussion to climate vulnerability in Indonesia, showing how integrated statistical, environmental, and geospatial data can reveal risks to fisheries and coastal livelihoods that often remain hidden in conventional indicators. Together, these articles highlight how statistical systems help societies identify problems that might otherwise remain unseen.

Learning through data

EDUCATION AND STATISTICAL LITERACY.

Francisca Ubilla and colleagues connect drought and water scarcity to real-world data analysis through the CR2 Climate Explorer, illustrating how secondary data can support inquiry into pressing environmental challenges. Reija Helenius and Elisa Falck reflect on the role of statistical literacy across all stages of life through the International Statistical Literacy Project. Together, these contributions demonstrate how data can serve as a bridge between everyday experiences, global challenges, and informed citizenship.

Statistics in action

COMMUNITIES, TECHNOLOGY, AND PUBLIC SERVICE.

Claire Kelling and colleagues reflect on community-engaged statistics and the importance of collaboration, participation, and lived experience in social research. The interview with Tony Labillois continues this discussion through themes of accessibility, official statistics, and public service. Complementing these perspectives, David Haughton and colleagues examine the design of computational workflows for human rights archives, highlighting how software infrastructure can support the responsible use of sensitive data. Across these contributions, statistical practice emerges not only as technical work, but also as a civic, communicative, and increasingly technological responsibility.

How can statistics help make hidden realities visible and support a more informed society?

Fabrizio Ruggeri
President, International Statistical Institute

Keep Learning, Keep Serving

What drives ISI and the people behind vISIon

In the interview in the second issue of vISIon, Tony Labillois mentions "energy, passion, curiosity." These words perfectly fit the work of the magazine's editorial board, and ISI is deeply grateful to all its members. vISIon has just begun its journey, and its spirit is already evident, as is that of ISI. In Tony's words, it is driven by a "very strong desire to serve, to give back, to transmit the values [we] have received, and to keep learning." The same principles underpin the articles in this issue: different problems, different approaches, but with the common desire to make a difference. This is what ISI aims to do, and vISIon is doing. Enjoy the read, and don't hesitate to leave a comment!

> *"Energy, passion, curiosity."*

Manuele Leonelli
Editor-in-Chief, vISIon: The ISI Magazine

Data for a Fairer World

On inequality and the statistician's responsibility

We live in a world of widening gaps. Between nations and within them, inequalities in income, health, education, and opportunity have deepened in recent years, accelerated by crises that exposed how unevenly risk and resilience are distributed. The United Nations Sustainable Development Goal 10 – Reduced Inequalities – reminds us that a fairer world is not only desirable but measurable. And where something is measurable, statisticians have a role to play.

That conviction runs through every article in this issue. One contribution reflects on race and official statistics: a reminder that the decision of what to measure, and what not to, is never neutral, and that data can both document injustice and, if handled carelessly, reinforce it. Another explores how community-engaged research can transform the way statisticians work on social issues, making the research process itself more equitable. A third looks at how AI is being put to work against modern slavery, one of the starkest expressions of inequality in the global economy. And through the lens of statistical literacy, we ask how we ensure that the tools of data science reach everyone, not just those already at the table.

None of these are purely technical questions. They are questions about values, about who counts, and about who is counted. Statistics sits at their heart because inequalities that go unmeasured go unaddressed. Our discipline does not just describe the world: it shapes what we are able to see, debate, and change.

I am also delighted to draw your attention to two events on the ISI horizon. The **5th ISI Regional Statistics Conference**, the first ever held in Europe, takes place in **Valletta, Malta, 3–5 June 2026**, hosted at the University of Malta. It is a rare opportunity to bring the ISI community together in the heart of the Mediterranean, and I encourage you warmly to join. And looking further ahead, mark your calendars for the **66th ISI World Statistics Congress**, coming to **Lusaka, Zambia, 11–15 July 2027**.

Enjoy the read, and as always, reach out, push back, and help us build this together.

> *"Inequalities that go unmeasured go unaddressed."*

AI against modern slavery (AIMS)

Building an open foundation for assessing corporate accountability at scale

Abstract. Modern slavery affects approximately 50 million people, many of them exploited through private-sector supply chains. Modern Slavery Acts in the United Kingdom, Australia and Canada now require thousands of companies to disclose their countermeasures each year, yet most of these statements remain largely unanalysed because governments, NGOs and investors lack the capacity to review them at scale. Project (AI Against Modern Slavery) builds the open research foundation that gap demands. It combines the world's largest annotated dataset of modern slavery statements, fine-tuned language models, explainable cross-jurisdictional review, model distillation for low-resource deployment, and a quality-assessment framework that moves beyond legal compliance. Together, these papers and artifacts form a transparent, deployable foundation for assessing corporate human rights compliance, in direct support of UN Sustainable Development Goal 8.7. The matching infrastructure on the policy side, machine-readable submission registries, standardised criteria and clear enforcement and supervision, has yet to be built by governments.

Adriana Eufrosina Bora
Adriana led Project AIMS at Mila, the Quebec Artificial Intelligence Institute, and conducted her doctoral research on the application of AI in reading and benchmarking modern slavery statements at Queensland University of Technology. Her work has appeared at leading machine learning venues, including ICLR and ACL, and has contributed to policy reviews of modern slavery legislation in Australia and beyond.

An estimated 50 million people live in modern slavery today, and approximately 86% of those in forced labour are exploited inside private-sector supply chains. For a decade, the legal mechanism designed to expose this, the Modern Slavery Acts (MSAs) in the UK, Australia and Canada, with New Zealand soon to follow, has quietly underdelivered. The problem is not simply that companies fail to disclose; many file year after year. The gap is on the receiving end: no stakeholder has the capacity to read the thousands of statements that arrive annually and turn them into insights that regulators, investors or advocates can act on.

The arithmetic is stark. The UK collects between 12,000 and 17,000 statements annually. Australia receives over 3,500. Canada, more than 6,000. Most statements run to ten or more pages. Even a well-resourced analyst could not work through one country's pile in a meaningful timeframe, let alone three. The transparency mandate exists; the data infrastructure to deliver on it does not. What cannot be assessed cannot be enforced; disclosure thus becomes a paperwork exercise, and laws meant to advance UN Sustainable Development Goal 8.7, ending forced labour and modern slavery, are kept from delivering on it.

Project AIMS (AI Against Modern Slavery) exists to close that gap from the research side. It builds the open foundation that scalable assessment will need: datasets, models, evaluation frameworks, governance protocols. The matching infrastructure on the policy side, machine-readable filing registries, standardised criteria, mandatory data formats, and clear enforcement and supervision, has to come from governments. Both halves are needed. Project AIMS is what the research community offers toward the first.

Building the foundation

Project AIMS is not a single tool. It is an open research foundation: datasets, models, frameworks and governance protocols built so that governments, NGOs, investors and researchers can assess corporate human rights compliance at a scale that matches the laws. The project has four main components built on academic papers: AIMS.au for annotated data, AIMSCheck for explainable compliance review, AIMSDistill for efficient deployment, and AIMS-QA for disclosure-quality assessment.

The project was prototyped by Adriana as part of her master's thesis in 2017 and was formally launched in 2020 by a partnership between The Future Society and Walk Free around a simple research question: Can machine learning accelerate the review of modern slavery statements? Five years later, after a second partnership between Mila and the Queensland University of Technology, the answer is yes, with the caveat that "yes" only holds if the AI is built carefully, transparently, and on data the field can trust.

> *"You cannot build trustworthy AI without trustworthy data. Every model trained on this dataset learns from human judgment about what compliance actually means."*

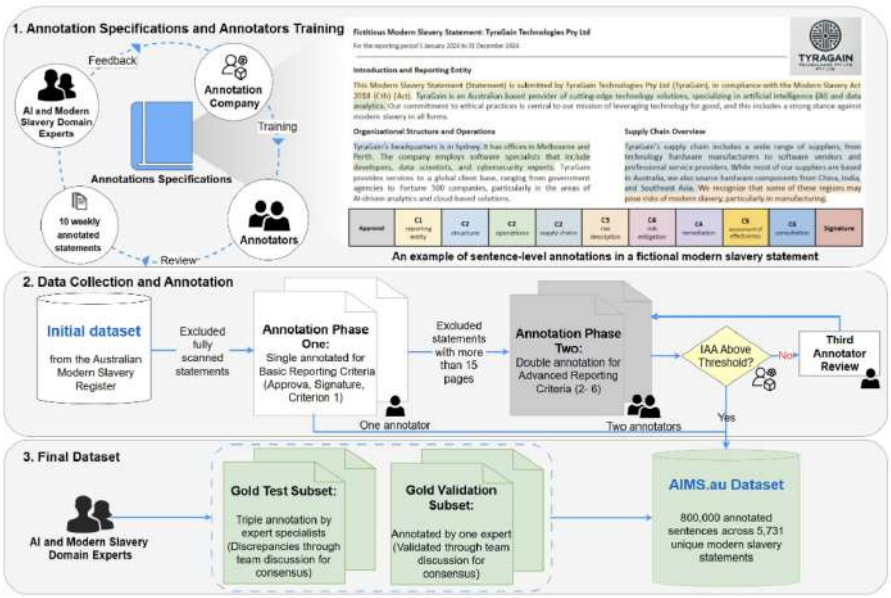

Figure 1. The AIMS.au annotation pipeline. Modern slavery statements published on the Australian Modern Slavery Register are converted from PDF to sentence-segmented text, then reviewed by trained annotators against eleven granular labelling questions derived from the Act's seven mandatory criteria. Statements falling below the inter-annotator agreement threshold are routed to a third annotator for adjudication.

The first contribution is AIMS.au (Figure 1) (Bora et al., 2025a). We collected 5,731 statements filed under Australia's Modern Slavery Act, covering more than 7,200 reporting entities across over twenty industrial sectors. A multidisciplinary team (machine learning researchers, business and human rights experts, modern slavery specialists) designed an annotation protocol that breaks the Act's seven mandatory criteria into eleven sentence-level labelling questions. Trained annotators then read each statement and marked which sentences addressed which criterion, including supply-chain risks, mitigation taken, consultation processes, effectiveness assessments, and remediation mechanisms.

Quality control was the slow, expensive part of the work, and the part that mattered most. Roughly 82% of the corpus was double-annotated, with statements falling below an inter- annotator agreement threshold routed to a third annotator for adjudication. Two "gold" sets of fifty statements each, jointly produced by three in-house experts, serve as benchmarks for any trustworthy model.

AIMS.au is not only a dataset. The same paper introduces a suite of fine-tuned language models, variants of DistilBERT, BERT and LLaMA, trained directly on the annotations. It also compares them with the alternative: zero-shot prompting of general-purpose LLMs, including GPT-4o. The result is a useful corrective. The smaller fine-tuned models beat the zero-shot LLMs, sometimes by wide margins. Performance also improves substantially when each sentence is evaluated with surrounding context (a window of roughly 100 words on each side), which mirrors how a human reviewer actually reads a statement.

All of it is open: the dataset, the annotation protocol, the fine-tuned model weights and the evaluation code, published on GitHub and Hugging Face under a permissive licence. Any researcher, regulator or NGO can audit and extend the work, and every subsequent paper in the AIMS series carries those fine-tuned models forward, retraining and extending them on new annotation data as the project moves into new jurisdictions and new questions. The dataset and the models do not belong to us. They belong to the ecosystem trying to end modern slavery.

Compliance review, with explanations

Working with this complex data made one requirement clear: explainability had to be built in from the start. A regulator who is shown a "compliant" or "not compliant" label by an AI model has every right to ask why. AIMSCheck (Bora et al., 2025b) is the framework we built to make that question answerable, with explainability as a foundational layer rather than something added later.

When the system assesses a statement, it does three things at once. It identifies the specific sentences that triggered each assessment. It uses SHAP (SHapley Additive exPlanations) values to highlight the words that most influenced each decision. And it labels the evidence status of each detected sentence: an action already implemented, a future commitment, or an explicit negative disclosure where the company states it has not acted.

As part of AIMSCheck (Figure 2), we tested whether models trained on Australian data could generalise to UK and Canadian statements. We built two additional annotated datasets, AIMS.uk and AIMS.ca, each with 50 official statements labelled against nine harmonised cross-jurisdictional criteria. Although the three Modern Slavery Acts overlap substantially, they differ in key respects: the UK recommends rather than mandates reporting criteria, Australia requires consultation with owned or controlled entities, and Canada focuses on child and forced labour, including remediation for income loss affecting vulnerable families.

> *"The transparency mandate exists. The data infrastructure to deliver on it does not."*

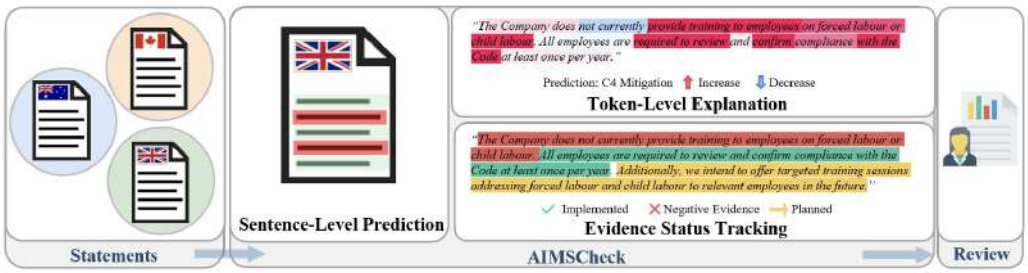

Figure 2. The AIMSCheck pipeline. Pre-processed statements are passed to a fine-tuned classifier that predicts sentence-level relevance against each reporting criterion; SHAP token attribution highlights the words driving each decision; and an evidence-status tracker labels detected sentences as describing an implemented action, a future commitment, or a negative disclosure. The three layers are surfaced together so that human reviewers can audit and confirm every prediction.

We then carried the fine-tuned models from AIMS.au forward and tested them on the new datasets, rather than starting from scratch for each jurisdiction. They generalised well: a 3-billion-parameter LLaMA model with 100 words of context achieved competitive macro-F1 across all three datasets, with only modest performance loss on UK and Canadian statements. Cross-jurisdictional AI assessment is feasible, but not by default. Each new country still demands real work. What changes is the starting point.

From research to practice

The practical reality is that the best research models are often the worst deployment candidates. Our top systems combined multiple fine-tuned language models with sophisticated post-processing, achieving state-of-the-art accuracy at a computational cost that governments and NGOs in resource-constrained settings simply cannot bear.

This is not just an optimisation problem. It is a justice problem. If compliance assessment can only be afforded by wealthy stakeholders, the resource-constrained ones, often the actors closest to the harm, are left behind.

AIMSDistill addresses that asymmetry (Bora et al., 2025c). Using knowledge distillation, we compressed what an ensemble of specialised teacher models had learned into a single 340-million-parameter student. The student runs roughly seven to eight times faster than the ensemble it replaces, uses about six times less energy, and matches performance across the AU, UK and CA test sets, all on the kind of standard 10-GB GPU available in resource-constrained deployments (Figure 3).

Beyond compliance, toward quality

Compliance and quality are different problems. A statement can address every mandatory criterion of the law without saying anything substantive: risks "have been assessed", remediation "has been established", effectiveness "is being monitored", all without naming a single supplier, a single risk, a single corrective action. AIMS.au and AIMSCheck are built to detect compliance. AIMS-QA, our most recent framework, is built to assess quality. (Bora et al., 2025d)

The starting point was a long-standing problem in the field. Academia and civil society have spent the last decade building their own benchmarking methodologies for what good disclosure should look like, beyond what the law mandates. Investor benchmarks, civil-society scorecards, NGO indices, sector-specific assessments. Each is a careful, manual scoring rubric. Together they cover hundreds of indicators that the law does not explicitly require but that anyone serious about supply-chain accountability cares about: living wage, supplier audits with worker engagement, training rates, supplier-tier transparency beyond tier one, grievance mechanisms, etc. The problem is that they are not comparable to each other and cannot be applied at scale.

AIMS-QA closes both gaps. We synthesised 13 publicly available benchmarking methodologies into a single taxonomy: 466 distinct metrics consolidated into 211 streamlined questions organised under 117 thematic topics, each explicitly mapped to the seven mandatory criteria of the Australian Act. Within that taxonomy, 37 questions appear in at least three independent methodologies; we treat

50 million

people live in modern slavery today

Approximately 86% of those in forced labour are exploited inside private-sector supply chains.

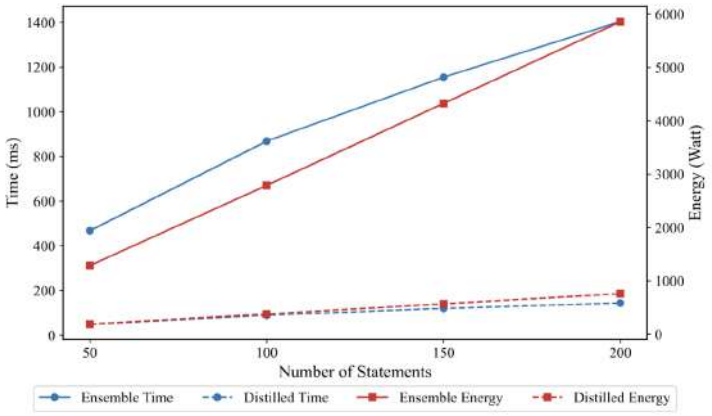

Figure 3. Efficiency comparison between the AIMSDistill student model (340M parameters) and the four-teacher ensemble (≈8B parameters) when classifying 200 modern slavery statements on a 10 GB GPU typical of resource-constrained deployments. The distilled student processes the same workload roughly seven times faster than the ensemble and uses about six times less energy, with no loss in macro-F1 across the AU, UK and CA test sets.

those as the consensus core, the indicators civil society and academia agree should be assessed regardless of stakeholder priorities.

The assessment itself runs through a two-agent pipeline (Figure 4). Agent 1 reuses the distilled student model ModernBERT classifier carried forward from AIMSDistill to identify the sentences relevant to a given question, then passes those sentences with their surrounding context, the question, and a chain-of-thought prompt to a large language model (we use Gemma-27B). The LLM is instructed to ground every answer strictly in the supplied evidence and to return a yes/no answer with explicit reasoning and evidence. When Agent 1 cannot find sufficient evidence, or when the LLM signals uncertainty, Agent 2 activates: a retrieval-augmented generation (RAG) watchdog that re-chunks the statement, uses a fine-tuned SentenceBERT retriever to find semantically relevant passages, and re-runs the question against the retrieved evidence.

The output to the human reviewer is never a bare yes or no. It is the answer, the model's reasoning, and the supporting evidence, presented together so the reviewer can verify or overrule. On initial evaluation, AIMS-QA achieves overall good accuracy against two independent domain experts yet some important errors were noticed. The errors are instructive: occasional hallucinations, compound questions that the model treats as a single test (and so dismisses partial evidence), and consistent failure to distinguish between concepts that look similar but differ legally, like minimum wage versus living wage, or a whistleblowing policy versus a worker-accessible grievance mechanism. These are not just engineering bugs. Some of them are the same distinctions that human reviewers struggle with, and the framework is designed to surface them rather than smooth them over.

A real-world example: the AIMS Hackathon

One way to test whether a research foundation is real is to see whether anyone outside the project can use it. In October 2025, the AIMS Hackathon, co-organised with Fundación Pasos Libres, Mila and QUT, convened the global AI community around the project's open datasets and models. Over two weeks, 227 participants from 38 countries formed 51 teams, with a final 23 teams submitting working solutions across four tracks: data mining and enrichment, model optimisation and explainability, stakeholder applications, and blue-sky innovation.

> *"The harder the data and the higher the stakes, responsible AI design stops being a nice-to-have and becomes a precondition for the work being trusted and used."*

Winning entries included a platform for generating structured statement summaries with explainable validation, accessibility-focused stakeholder applications, and large-scale data-extraction systems. More than fifty mentors from academia, civil society, government and industry ran workshops on legislation, due diligence, ethical AI design and NLP for social good. Post-event surveys reported that 95.9% of participants would recommend the hackathon and 87.7% wanted to participate again: useful evidence that openly released research outputs have the potential to seed durable communities of practice.

What's next

The map is expanding. New Zealand is preparing its own Modern Slavery Act. Germany, Norway and France already have their own mandatory human rights and environmental due diligence (mHRDD) laws on the books and many others are in development across Europe and beyond. Within a few years, the relevant global disclosure load will run not to tens of thousands of documents annually but to hundreds of thousands, each one addressing a wider range of obligations than today's modern slavery statements.

The analytical task gets harder too. Modern slavery legislation asks companies to describe what they did; the new generation of mHRDD laws asks regulators and stakeholders to verify it. That shifts the work beyond the statement itself, toward cross-referencing supplier registries, court filings, audit reports, NGO investigations, news cov-

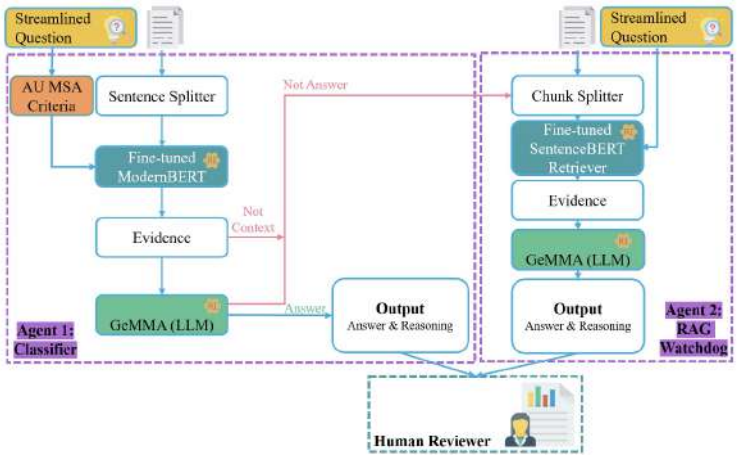

Figure 4. The AIMS-QA framework. A streamlined question and the source statement are passed to Agent 1, a sentence-level ModernBERT classifier that extracts the evidence relevant to the question and forwards it, with a chain-of-thought prompt, to a large language model for answer and reasoning. When Agent 1 cannot find sufficient evidence, Agent 2 activates: a RAG watchdog that re-chunks the statement, retrieves the most semantically relevant passages with a fine-tuned SentenceBERT retriever, and re-runs the question. The output to the human reviewer always includes the answer, the model's reasoning and the supporting evidence.

erage and trade data. Manual review of a single statement was already slow. Manual triangulation of a statement against a dozen external sources, repeated for hundreds of thousands of companies, is no longer slow. It is impossible. It is also more interesting AI work, on harder data, with larger consequences both for getting it right and for getting it wrong. The temptation, when the disclosure load is overwhelming and the deadline is real, is to throw whichever general-purpose LLM is at hand at the problem and hope for the best. The harder the data and the higher the stakes, the more responsible and ethical AI design (auditable evidence, calibrated uncertainty, human reviewers in the loop, openly reported errors) stops being a nice-to-have and becomes a precondition for the work being trusted at all

And the laws already in force are themselves moving. Australia is partway through the statutory review of its Modern Slavery Act, with a strengthened compliance regime under active consideration. The UK is debating reform of Section 54. Canada's guidance is being interpreted in real time as the first reporting cycles unfold. Methods built for one version of a law go stale when the law changes. AI assessment has to be designed to evolve with the legislation it is meant to monitor, not frozen at the point of publication.

This is where AIMS becomes a template rather than a finished product. The pattern (synthesise the expert benchmarks, ground a sentence-level dataset against the legislative criteria, fine-tune and distil for deployable accuracy, design explainability in from the start, validate with human reviewers in the loop) is reusable across every new disclosure regime. Building the next AI to support the monitoring of reporting under the mHRDD laws, or for New Zealand's legislation when it arrives, or for the next iteration of any of the Modern Slavery Acts, is no longer a project that begins from a blank page. The annotation protocols, the fine-tuned models, the explainability layers, and the evaluation methodology all carry over.

But research outputs are not products. AIMS today is a set of open datasets, papers, model weights and reference pipelines: the raw material a regulator or NGO would need to assess statements at scale, but not yet a tool a regulator or NGO can pick up and use without an in-house data or ML team. Closing that gap is the next major piece of work. A stakeholder-facing interface, queryable through plain language, that lets a compliance officer or a civil-society analyst load a statement, ask any of the 211 AIMS-QA questions or define their own, and receive the answer with its supporting evidence and the model's reasoning surfaced for review. The September 2025 AIMS Hackathon offered an early signal that this is achievable: 23 teams produced prototypes in two weeks of collaborative work. The science is in place. The community capacity is in place. What is missing is the sustained engineering investment that turns a fortnight of working prototypes into a tool a regulator can rely on year after year. The next phase, building those prototypes out into a sustained, maintained, civil-society-facing tool, is therefore more of a funding question. AI for human rights compliance occupies an awkward space: too applied for most academic funders, too public-interest for commercial ones. How that space gets bridged is one of the open questions for AI-for-good more broadly, and the answer for AIMS will shape the answer for many similar projects.

Years of building AIMS in close contact with regulators, civil society and academic colleagues has produced something the policy community can build on directly. Alongside the technical work, we have been collecting, refining and writing up a body of design lessons for legislators: what the next generation of human rights disclosure laws should look like if their output reports are to be assessable at scale, drawn from working hands-on with the disclosures the current generation has produced. Two halves of the architecture matter, and most laws today get one of them at most.

On the data side, disclosures should not arrive as free-prose PDFs; structured submission formats with required fields per criterion would make every downstream analysis cheaper and more reliable. Reporting entities need persistent identifiers that survive across years and jurisdictions, so a corporate group can be tracked through mergers, rebrands and sub-

sidiary reshuffies. Official registries should expose public APIs rather than search-only web interfaces. Mandatory criteria should be drafted with their assessability in mind, with each criterion written so it can be answered against evidence. Evaluation benchmarks should be funded as part of the legislation itself, not left as an afterthought for academics and NGOs to assemble. The other half of the architecture is enforcement. Disclosure laws without active supervision produce well-documented impunity rather than accountability; AI that surfaces non-compliance is only as useful as the body empowered to act on it. The pattern is visible in our data. In separate comparative work scoring twelve human rights due diligence laws against an "AI-readiness" framework, we found that monitorability and substantive ambition do not track each other: some of the most demanding due diligence regimes are among the least monitorable, while narrower laws with strong enforcement architecture score highest. The full analysis, including the AI-readiness framework and the scoring of all twelve laws, is set out in Bora et al. (2025e).

None of this requires governments to build their own AI. It requires them to build the foundation, on both halves of the architecture, on which any AI assessment worth trusting, public or private, would have to rest. Mandate the disclosure, design the infrastructure to read it, and resource the institutions that act on what is read.

The path forward

Modern slavery will not be solved by better data, and certainly not by AI alone. It will be solved by the long, unglamorous work of governments, businesses, workers, advocates and affected communities holding each other to account. Yet, that work needs tools. It needs to see patterns across thousands of statements and other datasets at once. It needs to know which companies fail to disclose, which make promises without evidence, and which genuinely engage with the workers and communities whose lives are most affected. Project AIMS is one piece of that toolkit: not the solution, but a foundation. The research is published. The datasets are open. The methods are proven. What remains is the work of turning artefacts into a usable interface, that interface into adoption by regulators and civil society, the lessons from years of research into the design of the next generation of laws, and that infrastructure into tools to help enforce and monitor the laws.

Acknowledgments: This work was supported by funds from the Commonwealth of Australia through the Research Training Program Scholarship, the Queensland University of Technology Centre for Data Science, the National Action Plan to Combat Modern Slavery 2020–25 Grants Program (administered by Australia's Attorney-General's Department) and Mila — Quebec Artificial Intelligence Institute. The author thanks Distinguished Professor Kerrie Mengersen (QUT), the co-authors and collaborators on the underlying papers, and collaborating organisations including The Future Society, Walk Free, WikiRate, Tasq.ai, and Fundación Pasos Libres.

Resources: Project page: mila.quebec/en/ai4humanity/applied-projects/ai-against-modern-slavery. Open datasets and code: https://github.com/mila-studios/ai4h_aims-au

References

Bora, A. E., St-Charles, P.-L., Bronzi, M., Tchango, A. F., Rousseau, B., and Mengersen, K. (2025a). AIMS.au: A dataset for the Analysis of Modern Slavery Countermeasures in Corporate Statements. Published in the proceedings of the International Conference on Learning Representations (ICLR) 2025.

Bora, A. E., Arodi, A., Zhang, D., Bannister, J., Bronzi, M., Tchango, A. F., Bashar, M. A., Nayak, R., and Mengersen, K. (2025b). AIMSCheck: Leveraging LLMs for AI-assisted review of modern slavery statements across jurisdictions. Published in the proceedings of the Association for Computational Linguistics (ACL) 2025.

Bora, A. E., Zhang, D., Bashar, M. A., Nayak, R., and Mengersen, K. (2025c). AIMSDistill: Distilling Knowledge from Specialised AI Teachers for Cross-Jurisdictional Compliance Analysis of Modern Slavery Statements. Submitted to Expert Systems with Applications.

Bora, A. E., Zhang, D., Thinyane, H., Bashar, M. A., Nayak, R., and Mengersen, K. (2025d). AIMS-QA: A Framework for Scaling Beyond Compliance the Assessment of Corporate Modern Slavery Statements with AI-driven Question Answering. Accepted at the Data for Policy Conference and submitted to the Data for Policy Journal.

Bora, A. E., Thinyane, H., and Mengersen, K. (2025e). Assessing the AI-Readiness of Human Rights Due Diligence Reporting Laws: Regulatory Design for Effective Reporting Compliance Monitoring. Submitted to Nature Humanities and Social Sciences Communications.

THE SCALE OF THE REVIEW PROBLEM

UK	12,000–17,000 modern slavery statements annually
Australia	More than 3,500 statements annually
Canada	More than 6,000 statements annually

Most statements exceed ten pages, making large-scale manual review impossible.

Beyond the data

How community-engaged statistics and data science transforms research on social issues

Abstract. Community-engaged research, in which community members participate throughout the research process, has the potential to improve the process and products of statistics and data science research on social justice issues. Drawing on perspectives from community members, students, staff of centers for community-engaged learning, and faculty, we illustrate how community engagement can shape statistics and data science research. We encourage statistics and data science educators and researchers to consider community engagement as an integral part of responsible and effective research on social justice issues.

Authors: Claire Kelling (Carleton College), Laurie Baker (Bates College, previously College of the Atlantic), Paige Bilich (Bates College), Carrie Diaz Eaton (Bates College, RIOS Institute), Tyler George (Cornell College), Morgan Kinney (Bates College), Jenny Mercado (Woonasquatucket River Watershed Council), Mary Parker (The People's Kitchen of San Luis Obispo), Emma Pederson (The Reinvestigation Workgroup), Angie Reed (Penobscot Nation), Emily Robinson (Cal Poly - San Luis Obispo), Anamika Sen (Bates College), Emily Seru (Carleton College), Carrie Slagle (Waypoint Services), Kai Zhang (Cornell University, previously Carleton College)

Introduction

Communities most affected by inequality are often excluded from or exploited by relevant research. People in the field of community-engaged statistics and data science are working to address this issue. Researchers should carefully consider when and how to ethically and sustainably include community members reflected in the data and/or affected by the results of their analysis. Community-engaged research emphasizes mutual benefit for both community members and researchers throughout the research process: communities gain access to statistical expertise and corresponding insights from data, while researchers gain contextual knowledge that a dataset alone can often miss. This approach has roots in public health and social science (Israel et al., 1998; El-Bassel et al., 2021), and is increasingly being considered within applied quantitative research (Olvera et al., 2024; Morton and Kelling, 2026). As statisticians and data scientists expand their roles in informing policy on issues of social justice and human rights, community-engaged frameworks offer a path toward research that is more informed, ethical, and impactful.

In this article, we center the voices of four groups involved in community-engaged statistics and data science research: community members, students, staff of centers for community-engaged learning, and faculty. In the following sections, we summarize the perspectives of coauthors from each group about why community engagement in statistics and data science research matters and how to approach it. Where appropriate, we provide attribution to specific coauthors' experiences through their initials.

Community members

Community-engaged research centers the perspectives of community members, including individuals and representatives from community partner organizations. These voices are crucial in developing research questions, conducting initial data cleaning and analysis, interpreting results, and disseminating findings. Research questions can come through intentional development or through more informal conversations between community members and researchers. EP is involved with the Reinvestigation Workgroup (RWG), a Minnesota-based volunteer team that investigates police violence and its causes. She reflects that at a meeting involving community members and researchers, an impacted family member shared that the officer who killed their brother was still working in that same police precinct. That moment sparked an idea - how can complaint data be used to model how internal police dynamics shape behavior leading to community complaints? This spurred months of collaborative work with the research team and RWG.

Throughout the data analysis process, community members' lived experiences also provide important context. MP, the president of the board for People's Kitchen of San Luis Obispo (SLO), noticed that student researchers saw the data differently after serving meals, where it's no longer just numbers, it's people. Community partners bring their lived experiences to these partnerships and researchers bring the statistical tools. In return, community members gain access to findings derived from data that are relevant to their lives and researchers leave with a deeper understanding of the community and research topic. These reciprocal relationships often extend beyond individual research projects and create opportunities for skill-building, collaboration, and sustained community engagement. Data skills workshops taught by external experts and hosted by LB (faculty) enabled AR and other community members to build new skills and meet new collaborators. Funding from the Henry David Thoreau Foundation supported one of LB's students to work with the Penobscot Nation Water Resources Program (PNWRP) as an intern for the summer.

> *"Years of meal counts were written on paper, but weren't really being used. They felt that the analysis makes a big difference when they are talking to donors, writing grants, or even planning for the week ahead."*

Table 1. Summary of community-engaged statistics and data science project teams considered throughout this paper, along with their community goals.

Project Topic/Title	Organizations (Coauthors) Involved	Community Goals
Community-Engaged Data Science Course: Monitoring Water Quality in the Penobscot River	Penobscot Nation Water Resources Program (AR), College of the Atlantic (LB)	Develop automated field and lab reports that identify and flag potential data quality issues to support the analysis of water samples on the Penobscot River for water quality monitoring.
Using Statistics to Understand Engagement Patterns	Waypoint Services (CS), Cornell College (TG)	Investigate the utilization and outcomes of a pilot program, the Homeless Diversion Program, which aims to reduce the risk of homelessness and break the cycle of recidivism by connecting justice-involved individuals with housing and essential stabilizing resources.
VECINA	Woonasquatucket River Watershed Council (JM), Bates College (CDE, PB), and other cooperating institutions	Visualize environmental and community information for neighborhood advocacy (VECINA) by gathering, distilling, analyzing, and making accessible data in a way that is useful for the communities that need it.
Transforming Data into Impact for Addressing Food Insecurity in San Luis Obispo, CA	The People's Kitchen of San Luis Obispo (MP), Cal Poly – San Luis Obispo (ER)	Digitize five years of paper meal records for People's Kitchen, build a sustainable data workflow using Google Suite tools, and develop an interactive dashboard to track and visualize trends in daily meal counts.
Analysis of Policing Data in Minneapolis	Reinvestigation Workgroup (EP), Carleton College (KZ, CK)	Characterize how internal police dynamics shape behavior leading to community complaints through the use of public complaint data.

Working with community partners also makes the research products and their dissemination more appropriate and meaningful. As co-facilitator of the Nuevas Voces project, JM explains that what started as a more technical tool became more accessible and bilingual because families shared their real experiences and needs. Simplifying complex data in this way resulted in tools that the Latino families in Providence can actually understand and use. We also note that complex data analysis and tools are not always the most appropriate methods to address the questions and needs of community members. A community-engaged research project has worked to understand pollution in the Penobscot River, which is seen as a relative to the Penobscot people. Contaminant levels in fish make it nearly impossible for the river to contribute safely to a traditional diet and Penobscot people to exercise their sustenance fishing rights. In collaboration with AR, students created tools for the PNWRP to summarize field and lab data for quality control review, a behind-the-scenes functionality that is foundational to their work and equally as valuable as flashy front ends. Connections with students, professors, and workshop instructors have continued to benefit the PNWRP directly and other Tribal environmental professionals through the Tribal Exchange Network Group.

Perspectives from community members have also emphasized the ability of statistical research to expand the capacity of community organizations that are often focused on day-to-day operations rather than data analysis. In her capacity with the non-profit Waypoint Services, CS emphasizes that these academic/community partnerships allow community organizations to continue focusing on critical, day-to-day services while academics help capture, analyze, and elevate data that might otherwise go underutilized. In work with her academic partner, CS notes that resulting insights strengthened the ability to demonstrate program impact, clearly connect services to outcomes, and build a compelling case for funding to support these efforts. Similarly, MP emphasizes that years of meal counts were written on paper, but weren't really being used. The organization leaders had a sense of what was happening, like things getting busier at certain times of the month at the People's Kitchen, but it was more of a feeling than something they could show. They felt that the analysis makes a big difference when they are talking to donors, writing grants, or even planning for the week ahead. In addition to impacts on funding and organizational planning, EP felt that their organization's collaborative approach with researchers, which was accessible to the public and shaped by community experience, became a practical tool for accountability and change.

These experiences from community partners suggest that community-engaged statistics and data science research contributes to a broader understanding of systemic challenges, ensuring research is both grounded in real-world experience and actionable for meaningful change.

Students

Taking part in community-engaged courses or research in statistics and data science has the potential to transform the student experience. Through participation in a community-engaged research experience at Carleton College, KZ noticed the local government data for his project quietly changing in impactful ways, with no press release or clear documentation announcing these changes. KZ found himself asking: If researchers with statistical training need weeks to make sense of a public dataset, how can community members hope to use it to hold institutions accountable? While data is often treated as a "given" in traditional statistics and data science courses, community-engaged research often involves important questions about the data-generation process, doc-

Figure 1. Photo of Laurie Baker's data science students from College of the Atlantic: Kristin Zunino, Linnea Goh, and Noelle Stringer learning about water sampling on the Penobscot river with Angie Reed and the late Jan Paul of the Penobscot Nation Water Resources Program.

umentation, and potential biases in the data sources. KZ noticed that working alongside community partners made clear that data access and literacy are often themselves the frontline of this work, not a preamble to it.

Through a community-engaged research experience at Bates College, PB similarly noticed that data on social and environmental issues is deeply shaped by human behavior and context through a collaboration with the Woonasquatucket River Watershed Council in Providence, Rhode Island. PB reflected that a particularly revealing moment came when the statistical analysis suggested that increased rainfall was associated with decreased flooding. Through conversations with community members, the group learned that during heavy rain events, people are less likely to submit reports, and some assume major floods are already documented. This insight was directly driven by community engagement.

Together, these examples highlight how community partnerships can fundamentally shift and expand student experiences in research and coursework through important community-driven contexts.

Community engagement staff

Staff at university centers for community-engaged learning are critical components of the process of implementing community-engaged research and teaching. They weave discipline-specific research and learning into community-engaged courses and research, aligning projects with community-identified priorities to achieve mutually beneficial learning and short- and long-term outcomes.

Community engagement staff often help facilitate relationships with community partners by encouraging best practices among faculty conducting community-engaged research. A staff member at Carleton College (ES) has identified that vital ingredients in effective community-scholar collaborations are humility, mutual trust, and respect. Specifically, when scholars enter communities with the attitude of a learner, with self-awareness that they don't know what they don't know, and when they are willing to take risks to see things from the community's perspective, an environment for deep collaboration is created. She emphasizes that some level of vulnerability is needed when faculty are asking others to share from a place of personal experience and real-life impact.

> *"If researchers with statistical training need weeks to make sense of a public dataset, how can community members hope to use it to hold institutions accountable?"*

Similarly, a staff member from Bates College (MK) has identified that one essential practice in community-engaged research is meeting the community where they are, including physically showing up but also making language and content approachable and useful to the relevant community. Academic jargon creates barriers and limits the potential for true collaboration and authentically applied research in community settings.

When community-engaged research on a social justice issue occurs within the context of a course, instructors can work with community engagement staff to design processes that ensure mutual benefit. From Bates College, MK illustrates this through an ongoing collaboration between three courses and a local youth-serving nonprofit, Rosati Leadership Academy, which teaches social-emotional skills through soccer. A multi-semester assessment project with Rosati resulted in the development of youth-appropriate survey measures to collect and analyze quantitative data on the social-emotional skill building work of the program. The results of this analysis will help the Rosati team persuasively tell their story in grant applications and fundraising meetings, with psychological theory-based evidence of their impact. The student researchers gained invaluable experience in translating theory into usable

measures and further translating those measures to a diverse young audience.

We strongly encourage statistics and data science researchers looking to pursue community engagement to work with and learn from the wealth of experience of staff at centers for community/civic engagement to design intentional and ethical projects.

Faculty

Community-engaged research involves many processes absent in traditional forms of research. For example, faculty must carefully consider how to build trust with community partners, and they cannot expect this trust to be immediate. In her partnership with the People's Kitchen of SLO, ER encountered initial skepticism about sharing data, but by showing up, serving meals, and meeting regularly, trust grew. This growing trust facilitated deeper engagement and analysis over time. In addition to physical presence, flexibility can be a powerful tool to build trust with community partners. CDE found that their group's willingness to pivot, during work with the Woonasquatucket River Watershed Council (WRWC), to an urgent real-time data question helped WRWC community members trust her group and their intentions. This pivot demonstrated that their work wasn't about the faculty's goals; it was about helping the community answer THEIR questions.

Flexibility is an essential part of community-engaged research. Conversations with community members over the course of the research process help to ensure that the work will be useful and impactful. However, these conversations can necessitate changes to analysis plans. What began as a data visualization effort by ER's team shifted into building something the organization could realistically maintain by creating a simple data entry system and dashboard that they have now been using for nearly three years. Similarly, although WRWC is an environmental organization, the deadline for school choice and lottery registration was impending during the first design sprint with CDE's research group. The pivot described earlier involved CDE's team temporarily refocusing their efforts on analysis of public educational outcomes data to answer a more pressing question: What is the best middle school for a student who is still learning to speak English? Flexibility in how results are disseminated became important in CK's collaboration on policing data analysis in Minneapolis. The form of the group's interactive data visualization was finalized after multiple rounds of feedback with community members, both individually and in larger groups, to ensure that the visualization answered questions of interest to the community. Flexibility of faculty is often necessitated in the classroom, but it is important to continue building these skills in research as well before developing community partnerships.

As discussed by student researchers, community partnerships can clarify findings from data analysis. Through ongoing conversations with community partners, ER's team learned the "why" behind the meal count trends. Spikes often aligned with benefits running out at the end of the month or not arriving until a few days into the next. This kind of insight reflects a broader emphasis shared across many courses on understanding data at its source. LB found that by engaging directly in the data collection process and working closely with community partners, students are better able to see both the origins of the data and its limitations, leading to more informed analyses and a deeper appreciation of the partner's work and the challenges they seek to address. Existing research, including theoretical findings, can also be translated into impact through partnerships with community organizations. Through product research for a citizenship loan that a local credit union was planning to launch, students in a Finance & Society class offered by AS were able to build on existing theoretical research to study the prohibitive costs of naturalization.

From the early stages of community-engaged work, faculty should be thinking about the sustainability of proposed solutions or research directions after individual projects or courses. TG has acknowledged that a single analysis might not be enough. While planning future initiatives, he asks how he can help his community partner continue to assess their programs after his class ends. Similarly, ER came to realize through her community-engaged research that a technically sound solution must also be usable long term. Identifying pathways for work to continue beyond the course, such as through follow-on student involvement in courses, independent research, or summer internships (as in LB and AR's research), can help sustain momentum and deepen impact.

These faculty illustrate that with some flexibility, community-engagement can transform research into a relationship-driven process that yields sustainable and deeply contextualized solutions.

Conclusion

Together, we offer the following advice and takeaways. Kelling et al. (2026) provide additional tips for implementing community-engaged statistics and data science courses, especially for forming learning goals and the planning stages.

- Start planning early and regularly involve, communicate with, and learn from community partners and staff from centers for community-engaged learning.
- Develop mechanisms that will help students understand the data from the community partner's perspectives.
- Include partners in the planning of data skills workshops that are relevant to their project needs and invite them to participate.
- Allow for community priorities to shift mid-project and pivot accordingly in order to deepen trust and impact.
- Explore routes for continuing projects with community partners beyond the course, such as securing funding for summer internships.
- When possible, lean towards simpler, sustainable solutions that partners can continue to use and maintain without needing specialized skills or training.

We hope this article will encourage researchers and educators to involve community members in their research processes when studying social justice issues. The perspectives of community members can enhance the impact of research through intentional structures

that motivate research questions, design, result interpretation, and dissemination efforts. Being open and willing to center research on community partnerships requires researchers and funders to step outside perfectly made plans and organized structures, but making a difference alongside our communities through statistics and data science research is the best reward.

> *"When scholars enter communities with the attitude of a learner, with self-awareness that they don't know what they don't know, and when they are willing to take risks to see things from the community's perspective, an environment for deep collaboration is created."*

Acknowledgements

We gratefully celebrate and pay tribute to Jan Paul, a strong Penobscot woman who spent 28 years protecting her sister the Penobscot River and a key member of the research collaboration between the Water Resources Program and College of the Atlantic. Hear more about her work as part of the Penobscot Nation to protect the water in this Posit video and learn more about her life at her memorial page. The voices of community members, students, staff, and faculty are all critical in shaping the community-engaged research discussed in this article; we thank all of the participants in the projects described in this paper.

References

El-Bassel, N., Gilbert, L., Hunt, T., Wu, E., Oga, E. A., Mukherjee, T. I., Campbell, A. N., Sabounchi, N., Gutnick, D., Kerner, R., Venner, K. L., Lounsbury, D., Huang, T. T., and Rapkin, B. (2021). Using community engagement to implement evidence-based practices for opioid use disorder: A data-driven paradigm & systems science approach. *Drug and Alcohol Dependence*, 222:108675.

Israel, B. A., Schulz, A. J., Parker, E. A., and Becker, A. B. (1998). Review of community-based research: assessing partnership approaches to improve public health. *Annual Review of Public Health*, 19(1):173–202.

Kelling, C., Baker, L., George, T., and Robinson, E. (2026). Tips for developing and teaching community-engaged statistics and data science courses. In *Proceedings of the Twelfth International Conference on Teaching Statistics*, Brisbane, Australia.

Morton, C. and Kelling, C. (2026). Towards community-based participatory statistics research using secondary data. *The American Statistician*, ((online)):1–24.

Olvera, R. G., Plagens, C., Ellison, S., Klingler, K., Kuntz, A. K., and Chase, R. P. (2024). Community-engaged data science (CEDS): A case study of working with communities to use data to inform change. *Journal of Community Health*, 49(6):1062–1072.

Using secondary data to engage learners in researching social issues

The case of the CR2 Climate Explorer and water scarcity

Abstract. Drought and water scarcity are phenomena that are often confused, yet they have different causes: drought has a natural origin, while water scarcity is caused by water use. In this article, we present a proposal to address these issues through the analysis of secondary data from the CR2 Climate Explorer. By proposing different types of research questions and following statistical research processes, we illustrate how the problem of drought and water scarcity could be addressed using data from the city of Petorca (Chile). We believe that this type of activity fosters student engagement and action regarding socio-environmental issues affecting different communities, allowing them to reflect on how data serves as a source of information for decision-making.

Francisca Ubilla
Assistant Professor at Pontificia Universidad Católica de Chile. Researches statistics education and civic statistics in teacher training.

Valentina Giaconi
Associate Professor at Universidad de O'Higgins. Holds doctorates in Mathematical Modelling and Education Sciences.

Raúl Valenzuela
Assistant Professor at Universidad de O'Higgins. Specialist in extreme hydrometeorological events, atmospheric rivers and radar meteorology.

Felipe Ruz
Associate Professor in the Department of Statistics at the Pontificia Universidad Católica de Valparaíso. Researches statistical and civic literacy in maths education.

Introduction

Education has a key role in forming citizens, present and future, who can make sense of complex social phenomena, with statistics offering a powerful interdisciplinary lens for such development (Engel et al., 2021). The analysis of different phenomena using data should occur through the connection between global and local issues; that is, we should understand the phenomena that affect us globally through the critical interpretation of data, while simultaneously reflecting on and gathering local evidence regarding how these issues impact the places we inhabit daily.

One of the phenomena that has been affecting a large part of the population for several years is water scarcity and drought, phenomena that appear similar, but are fundamentally different. "Drought" is a natural phenomenon caused by large-scale climate variability and cannot be prevented through local water management. "Water scarcity" refers to the unsustainable long-term use of water resources, which water managers can influence (Van Loon & Van Lanen, 2013). This issue also aligns with Goal 6 of the United Nations' Sustainable Development Goals (SDGs), which calls to "Ensure availability and sustainable management of water and sanitation for all." Water also cuts across nearly every other SDG as well. This makes water an ideal topic for engaging learners in that is related to both global and local issues. Based on this, the Center for Climate and Resilience Research (CR2) offers the opportunity to analyze the phenomenon of drought using a database derived from various weather stations belonging to the observation networks of the General Water Directorate (DGA), the Chilean Meteorological Directorate (DMC), and Chile's National Agroclimatic Network (RAN, Agromet).

The availability of water data in Chile has allowed us to develop activities based on the statistical problem-solving process proposed in GAISE II (Bargagliotti et al., 2021), which aims to encourage students to formulate statistical research questions based on secondary data. However, based on our experience and what the liter-

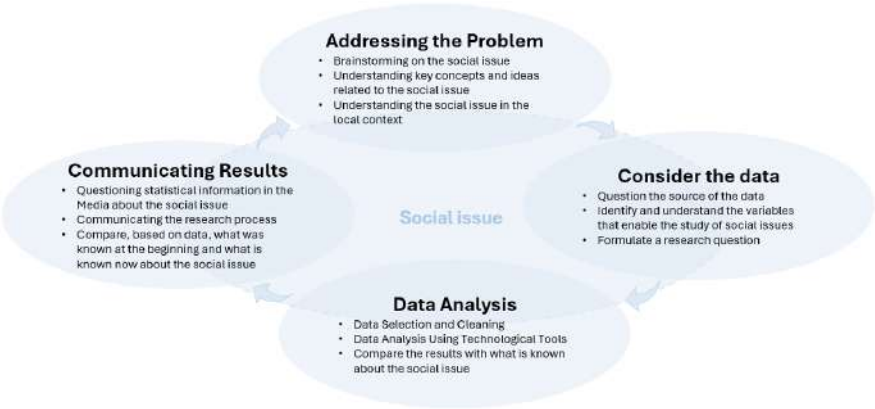

Figure 1. Statistical research process based on real data on social issues

ature reports (Ubilla and Gorgorió, 2023), formulating statistical questions is a complex process, which can be further complicated if students must generate such questions from an existing dataset. For this reason, we present ideas on how to address the issue of drought using a database by following the statistical research process, showcasing three examples of research questions that adhere to the guidelines by Arnold and Franklin (2021) on what makes a good statistical question.

How can we understand social issues through data?

Based on the GAISE II proposed statistical research process (Bargagliotti et al., 2021), in Figure 1, we outline some aspects of each phase, keeping in mind that the topic at hand concerns a social issue using secondary data.

First, we believe that when *approaching this issue*, it is necessary to identify what students already know about it and then introduce key concepts and ideas that will help them understand the phenomenon. Next, students are expected to learn how this issue affects their local community and what various institutions have done to address it. In this phase, they are integrating information about the social issue that stems from their own beliefs and experiences with theoretical, global, and local information on the issue. During the *consideration of data*, students are expected not only to question the source of the data but also to be able to identify the variables and nature of the data that help them understand the issue. The link between the previous stage and this one is established through the formulation of a research question. Because databases on social issues tend to be extensive and complex, during the *data analysis* phase, students are expected to clean the data so that, after analyzing it, they can compare the results with what they initially knew about the issue. Finally, it is hoped that, based on the results obtained, students will be able to compare them with information found in the media and integrate what they knew initially with what they observed from the data, and ultimately *communicate their process* to an audience they deem relevant.

> *"Education has a key role here, in terms of forming citizens, present and future, who can make sense of complex social phenomena, with statistics offering a powerful interdisciplinary lens."*

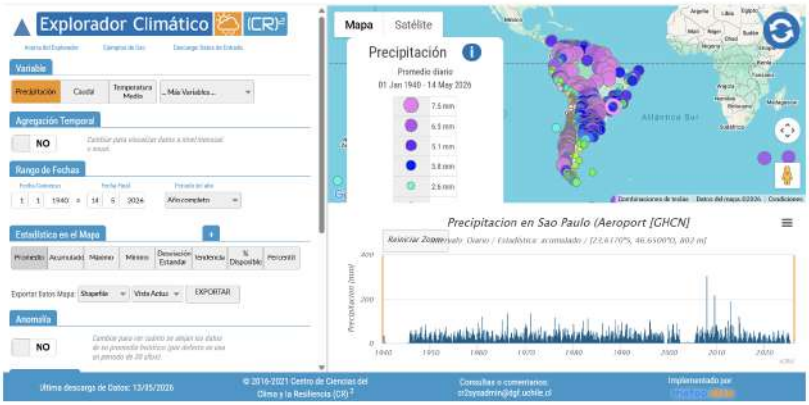

Figure 2. Home page of the CR2 Climate Explorer (https://explorador.cr2.cl/)

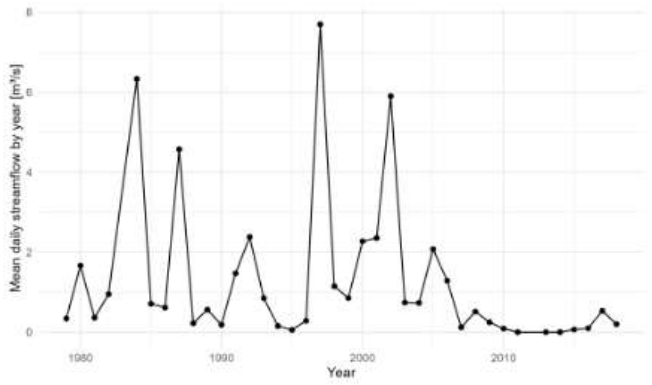

Figure 3. Annual mean daily streamflow of the Petorca River. The figure shows the yearly average of daily streamflow values, expressed in cubic meters per second [m³/s].

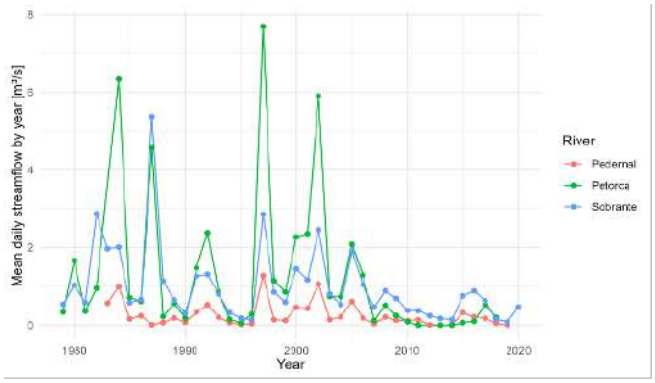

Figure 4. Annual mean daily streamflow by river. Values are expressed in cubic meters per second [m³/s].

What data can we find in the CR2 Climate Explorer?

The Climate Explorer from the Center for Climate and Resilience Research (CR2) compiles observational records of streamflow, temperature, and precipitation from climate and hydrological stations. Streamflow is expressed in cubic meters per second (m³/s), temperature in degrees Celsius (°C), including maximum, minimum, and average temperatures, and precipitation is the daily accumulation in millimeters (mm). The platform allows users to filter data by date range and time scale, view it in map or time series format, and export it for analysis. It contains data from 669 weather stations in Chile, South America, and Antarctica that provide temperature data; 809 weather stations in Chile alone that provide flow rate data; and 1,241 weather stations in Chile, South America, and Antarctica that provide precipitation data, that have been collected over the years. Figure 2 shows the Climate Explorer interface, where each point corresponds to a data station and at the bottom, you can view a time series of precipitation data. Click the following link to access the CR2 data viewer: https://explorador.cr2.cl/

In this article, we aim to demonstrate how to use local CR2 data to answer simple research questions, thereby illustrating how to work with this data in school settings or introductory statistics courses for undergraduate programs.

With this in mind, we decided to study the Petorca River due to the documented problems regarding access to and availability of water in this region. In particular, the National Institute of Human Rights (n.d.) describes Petorca as a case of water crisis linked to the depletion of surface water and restrictions on the use of groundwater.

In the CR2 Climate Explorer, we identified one station corresponding to the Petorca River and two stations associated with the Sobrante and Pedernal Rivers, both of which flow into the Petorca River. Since the issue under analysis is directly linked to water availability in the river system, we used streamflow and precipitation as our variables.

The analysis will use the data available in the three selected time series, covering the period 1979–2018. Although the original records are daily, annual averages for streamflow and annual totals for precipitation will be used to

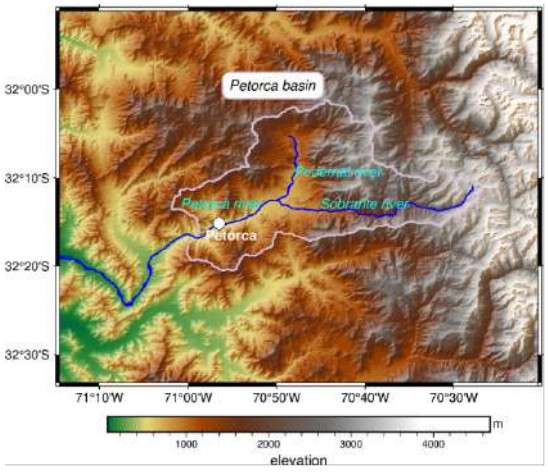

Figure 5. Topographic map showing the location of the three rivers in the mountainous landscape.

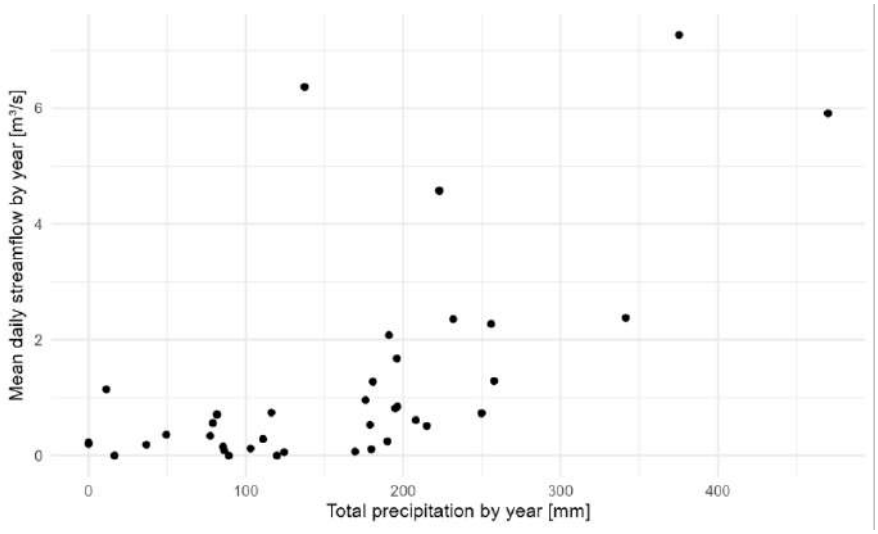

Figure 6. Relationship between annual precipitation and annual mean daily streamflow in the Petorca River.

simplify the analysis. After selecting the data, we downloaded the database and used the R software, version 4.5.3 to analyze it.

What does the data tell us about the drought in Petorca?

Our goal is to illustrate the approach and methods for addressing statistical research questions. To do this, it is essential to define the focus of the inquiry. In this proposal, we present three types of research questions: descriptive, comparative, and relational. The final question is presented along with the consolidated analyses; however, students are not expected to arrive directly at that final question or those results. The idea is to arrive at these types of questions through a dialogic process with teachers and peers.

> *"We should understand the phenomena that affect us globally through the critical interpretation of data, while simultaneously reflecting on and gathering local evidence regarding how these issues impact the places we inhabit daily."*

Summary investigative question. Descriptive or summary investigative questions refer to questions that characterize how a variable behaves within a group (Arnold & Franklin, 2021). In this context, such a question would begin by considering whether we are interested in describing a specific variable in one or more specific sectors. Considering the selected issue, a relevant initial question is How has the streamflow of the Petorca River behaved over the years? For each year, we calculated the mean daily streamflow in order to obtain simpler data. We plotted this variable against the year of the data, generating the graph shown in Figure 3.

We can observe a considerable variability between 1978 and 2006, followed by a very strong downward trend, revealing an important issue. This trend may be due to broader climatic phenomena or to human management of water resources.

Comparison investigative question. Comparison investigative questions refer to comparing how a variable behaves across several groups (Arnold & Franklin, 2021). To define them, it is necessary to identify the groups of data that we are interested in comparing. In the previous example, the issue affecting Petorca was made evident; however, it is not clear whether this pattern is due to broader climatic phenomena or to water management. To explore this further, we focused on the Pedernal and Sobrante rivers, which flow into the Petorca River. These rivers originate closer to the Andes Mountains and farther from the valley and productive activities, so their behavior could be more strongly associated with climatic phenomena than with direct human intervention (Figure 5).

We therefore asked: *Does annual mean daily streamflow tend to differ among the Petorca, Pedernal, and Sobrante rivers?* To address this question, we generated a graph that shows the yearly average of daily streamflow values for each river, allowing comparison of streamflow trends over time (Figure 4).

We observe that, until 2006, the Petorca and Sobrante rivers show considerable variability, whereas the Pedernal River consistently has the lowest streamflow. After 2006, the three rivers show a tendency toward lower streamflow, which suggests the influence of a climatic drought phenomenon affecting all of them. However, until 2006, the Petorca River tends to show higher streamflow than the other rivers; after this year, its streamflow is lower

than that of the Sobrante River and similar to that of the Pedernal River. This latter phenomenon suggests an effect of human action, in addition to broader climatic factors.

Association investigative question. Association investigative questions seek to explore whether there is an association between two variables (Arnold & Franklin, 2021). Continuing with the previous account, we considered that another way to understand whether the streamflow of the Petorca River has changed due to the megadrought or to human management is to compare streamflow with precipitation. Precipitation is less directly controllable by human action. Therefore, we asked: *Is there an association between annual precipitation and annual mean daily streamflow in the Petorca River?* To address this question, we plotted yearly precipitation values against yearly average daily streamflow values for the Petorca sector (Figure 6). We also calculated the Pearson correlation between both variables, which was $r = 0.70$, $p = 0.00$. The graph and the correlation show a strong but imperfect relationship between precipitation and streamflow — one that is not only non-linear in nature, but may even suggest an exponential pattern — leaving considerable room for other factors, including human action, to influence the river's behavior.

Implications

In this article, we present an example of data analysis that school students can undertake to answer research questions related to the issue of drought using the CR2 Climate Explorer (n.d.). The initial aim of this proposal was to address a common misunderstanding among the general public regarding the phenomena of drought and water scarcity. By analyzing secondary data and posing questions that seek to address the issue from different perspectives, we have characterized the phenomenon of drought in Petorca, where the interpretation of graphs has allowed us to question the climatic nature of this phenomenon and bring human intervention into play in this water crisis. The water crisis in Petorca has been studied, and scientists specializing in climate science indicate that the crisis is due to human action and not solely to a megadrought (Álamos et al., 2023; Muñoz et al., 2020). In this vein, this activity is valuable because, through simple data analysis, it is possible to reflect on and question human action in the context of the climate crisis. Furthermore, the CR2 Climate Explorer platform offers an opportunity for the Latin American region, as it provides access to data from various weather stations. In addition, we consider it essential to encourage students to develop data cleaning skills (Wild & Pfannkuch, 1999), since databases on social issues often contain large amounts of data and variables that sometimes need to be processed in order to answer research questions. All of this is done so that students can then use software that allows them to decide which analyses are needed to answer the questions they pose, as well as to create a space for reflection and dialogue between their experiences regarding the issue and what the data reveal about it.

Acknowledgements

F.U. thanks the support of ANID Fondecyt Iniciación 11250974, V.G thanks the support of ANID/PIA/Basal Funds for Centers of Excellence Center for Mathematical Modeling (CMM) FB210005., R.V. thanks the support of ANID Fondecyt Iniciación 11230184 and the Center for Climate and Resilience Research (CR2), based at the University of Chile and funded through the Actividades de Interés Nacional (AIN) program.

References

Álamos, N., Alvarez-Garreton, C., Muñoz, A., and González-Reyes, A. (2023). The influence of human activities on streamflow reductions during the megadrought in central Chile. *Hydrology and Earth System Sciences Discussions*, 2023:1–28.

Arnold, P. and Franklin, C. (2021). What makes a good statistical question? *Journal of Statistics and Data Science Education*, 29(1):122–130.

Centro de Ciencia del Clima y la Resiliencia (CR2) (2025). Inicio. https://www.cr2.cl.

Instituto Nacional de Derechos Humanos (2025). Petorca, historia de una crisis hídrica. rehuir el olvido. https://shorturl.at/OJJr2.

Muñoz, A. A., Klock-Barría, K., Alvarez-Garreton, C., Aguilera-Betti, I., González-Reyes, A., Lastra, J. A., and LeQuesne, C. (2020). Water crisis in Petorca Basin, Chile: The combined effects of a megadrought and water management. *Water*, 12(3):648.

Ubilla, F. M. and Gorgorió, N. (2023). Using Eurostat data to teach statistics to prospective primary teachers: on how the context of the task may promote their social awareness. *Journal of Mathematics Teacher Education*, 28:655–680.

Wild, C. J. and Pfannkuch, M. (1999). Statistical thinking in empirical enquiry. *International Statistical Review*, 67(3):223–248.

The International Statistical Literacy Project

Putting the Focus on Statistical Literacy

Written by: Reija Helenius (ISLP Director) and Elisa Falck (ISLP Project Coordinator)

The International Statistical Literacy Project (ISLP) is a nonprofit project operating within the ISI and the International Association for Statistical Education (IASE), dedicated to promoting statistical literacy worldwide in all domains and stages of life. This article outlines the history, activities, challenges, and future directions of the project, drawing on articles written together with Pedro Campos, Adriana D'Amelio, and Steve MacFeely.

Statistical literacy: what is it?

Statistical literacy is the ability to understand, interpret, and critically evaluate data and statistical information in everyday life. In a world where we are constantly surrounded by numbers (for example, news, social media, and public discussion), it is an essential skill for making sense of the information we encounter. The International Statistical Literacy Project (ISLP) is aware of the need to promote statistical literacy in all domains and stages of life. This is especially relevant concerning sustainable development goals, for instance with regard to environmental concerns and education.

Being statistically literate means more than just reading charts or knowing basic concepts. It involves asking questions, recognising possible biases, and understanding how data are collected and presented. It also helps people judge the reliability of claims and distinguish between strong evidence and misleading information. In democratic societies, the ability to think critically about information is inseparable from the right to access trustworthy sources, empowering individuals to make autonomous decisions and participate meaningfully in public life.

Ultimately, statistical literacy empowers individuals to make informed decisions in their daily lives, from personal choices to societal issues. It supports active participation in communities and strengthens our ability to engage with an increasingly data-driven world. Some new challenges have arisen as AI-generated information enters everyday decision-making. People need to be equipped with the skills to determine whether data is trustworthy, especially data produced by AI systems.

The history of ISLP and its current administration

The ISLP is a nonprofit project that operates within the International Association for Statistical Education (IASE) and the International Statistical Institute (ISI). The mission of the ISLP is to support and promote statistical literacy activities around the world and in all areas of life, especially in less developed countries. The project was started in 1994, when an ISI committee called World Numeracy Program (WNP) was established to stimulate the spread of quantitative skills all around the world. The person who started it all was Luigi Biggeri from Italy. Under the lead from Biggeri, a group of action proposals intended to support the project was compiled. In 2000, the ISI President invited this programme to come under the umbrella of the International Association of Statistical Education (IASE). Carol Joyce Blumberg (USA) took over as chair in 2001. After this, Juana Sánchez (USA) was appointed as Director in October 2006. Reija Helenius (Finland) has been leading the project since 2010.

Sanchez continued Blumberg's work in developing the ISLP network with volunteers. Both wanted to also bring national statistical offices more actively into the work of the ISLP. The statistical literacy competition for young people was launched. In 2007, a pilot competition was organised in Portugal and in 2008 to 2009, a global competition that reached its climax in the international final during the ISI Biennial Session in Durban, South Africa. The second competition, the Best Cooperative Project Award was also launched.

The worldwide country coordinator network was formed in 2010. At the same time, the format of the ISLP International Poster Competition was made, and the first competition launched. These networking efforts have now led to a group of 184 country coordinators in 97 countries around the world. This project would not be possible without the dedication of our volunteers—particularly the country coordinators. They organize events within their respective countries, including webinars and poster competitions.

Volunteer work has achieved much in different countries and for the organisation of the project as well. From the perspective of development of the cooperation network, it is important to tighten collaboration between different actors inside each country. Such actors are national statistical institutes, statistical societies, universities and educational institutions. Collaboration with the media and the library network is also important. We also warmly invite you, the reader of this article, to visit the ISLP site (https://iase-web.org/country-coordinators) and see if there is a national ISLP Country Coordinator in your country. We encourage you to contact them and support their statistical literacy promotion work, or become a Country Coordinator yourself.

The ISLP has two committees: the Executive team which includes Reija Helenius (Finland), Pedro Campos (Portugal), Irena Ograjensek (Slovenia), Adriana D'Amelio (Argentina) and Saleha Habibullah (Pakistan), and the ISLP Advisory Board, chaired by Steve MacFeely (Ireland). The ISLP Advisory Board and Executive Committee have been nominated to serve until further notice.

The activities of the ISLP

The ISLP's biggest current project is the International Poster Competition, in which young people create a statistical poster. The format is intentionally designed so that the competition can be organized anywhere in the world.

We share our experiences and activities in different countries via the ISLP newsletter. The ISLP website also serves as a source of knowledge. The project keeps growing and adapting to the changing world.

Below we will briefly describe the activities of ISLP.

International Poster Competition. The International Poster Competition began in 2010 with the participation of several countries representing six regions of the world, as can be seen in Table 1. The competition has evolved over the years. In 2018, the idea arose from the ISLP to launch the partial Latin American Poster Competition with the need to increase the participation of countries in South America and the Caribbean. In 2020, the undergraduate category was added. In 2024, the competition was extended to primary school children aged 9 to 12. The competition is currently in its 11th edition, with submissions open until March 2027. The competition has four categories: primary schools (ages 9–12), lower secondary schools, upper secondary schools, and bachelor-level university students. In the competition, students conduct a statistical study and based on the study, design a statistical poster in teams of 1–5 students about a topic of their choosing. The competition facilitates young people's learning of statistical thinking, from thinking of a research problem to choosing a study method, data collection and analysis, to presenting data and making conclusions. The latest competition in 2024–2025 had 22,878 participants from 36 countries (see Table 1).

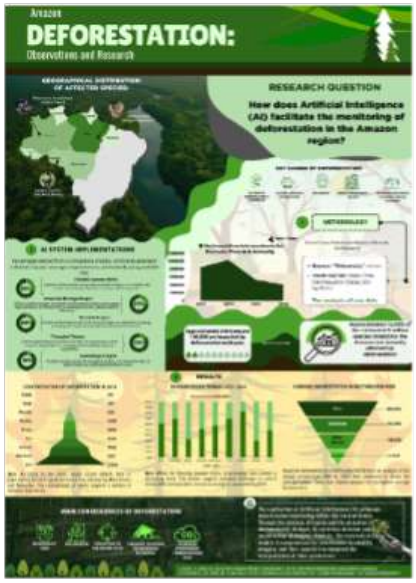

Figure 1. Example of a statistical poster. First prize winners in the Upper Secondary category, from Ecuador, in the International Poster Competition 2024-2025.

Creating a poster is a pedagogical tool in teaching periods for statistics. In many countries, this is a new teaching method that also fosters the creativity and independent thinking of students. As the poster competition participants work in teams of one to five people, the competition supports team-based learning and brings out the students' strengths in different areas. One student can be mathematically talented, another visually, and another can have a strong command of the phenomenon studied. Creating posters also supports phenomenon-based learning, and teaching based on preparing posters can be used as a communal project with other disciplines (e.g., mathematics, history, art, information technology, biology). A characteristic of the competition is also promoting critical literacy, particularly a command and understanding of numerical information. These skills are needed in both work and personal life.

Teachers who have participated in the competition have given positive feedback. Young people have described the statistical study process as meaningful because they can study topics in their environment, for example. The competition also inspires youth who have not yet embraced statistics.

Best Cooperative Award. The ISLP organises the Best Cooperative Award competition for the best statistical practices. The prize is given every two years to a statistical literacy project recognised as outstanding, innovative, and influential and reaching a broad segment of the general public. The projects that have received the award are the fruit of the cooperation of different types of institutions (national statistical offices, schools, statistical societies, media, libraries, etc.).

Submitted projects should possess specific attributes, including being current and having the potential for future relevance. They should also provide open access to project resources. The projects should educate on statistical theory and data analysis concepts, illustrating their application in generating information about countries and societies. Furthermore, the contents of these projects should be pedagogically sound and appropriate for a general audience. Collaboration between institutions that do not typically work closely together, such as projects fostering cooperation between different entities, is particularly encouraged.

The winner from 2024-2025 was StatBel Academy/Statbel Junior. "With Statbel Academy, Statbel is reaching out to education, and wants to support teachers in their approach. At the same time, future projects are also possible for other target groups. By continuing to focus on statistical literacy, including among children and young people, Statbel wants to contribute positively to a better-informed society in which citizens build up the necessary knowledge to deal critically with data and figures, and so make well-informed decisions."

The six previous winning entries include What's Going On In This Graph?", an excellent free weekly educational feature co-created by the American Statistical Association and the New York Times Learning Network, TREND - from Russia, a school competition in statistics conducted by the Russian Association of Statisticians; LeME from Brazil, a Statistical Multimedia Literacy project that involves collaboration between universities, other schools, communities, national societies and the Brazilian Ministry of Education. In 2017, the National Schools' Poster Competition from Australia won with collaboration among universities, national societies, industry professionals, schools, media experts, and sponsors, with additional support from the Australian Govern-

Table 1. Countries entering the ISLP International Poster Competition by region, 2012–2025.

Region	12–13	14–15	16–17	18–19	20–21	22–23	24–25
Africa	2	8	11	6	2	3	6
Asia	10	6	7	6	11	8	8
Europe	13	12	11	12	10	9	9
N. & C. America	2	3	2	3	3	3	4
South America	2	3	4	7	7	7	8
Oceania	2	2	2	2	0	1	1
Total	31	34	37	36	33	31	36

Source: ISLP website

ment's Department of Education and Training. In 2015, the winner was Exploristica, which was a cooperation between the NSO and national societies. You can see earlier winners, from the years 2009-2014, in the ISLP webpage.

International Day of Statistical Literacy (IDSL). The International Day of Statistical Literacy (IDSL) is a free, global online event that celebrates the importance of statistical and data literacy. The main event is held annually, is conducted online, and is open to all participants at no cost. IDSL brings together educators, practitioners, and enthusiasts from around the world to exchange ideas, share experiences, and showcase innovations in statistical literacy. Each year, the programme features presentations and keynote addresses by leading experts in statistics and education, lightning talks highlighting innovative projects and teaching practices, as well as interactive discussions and opportunities to connect with colleagues worldwide.

The next IDSL will be held on 17 November 2026. Mark your calendars! More information will be available soon on the ISLP website.

Other activities. The ISLP disseminates information about statistical literacy through the ISLP Newsletter, by organising webinars about different topics, sessions, workshops in statistical conferences, and events in the educational sector. It also has its website (https://iase-web.org/islp), and it supports country coordinators with advice and instructions. The project is involved in initiatives such as collecting statistical literacy information sources with GIST (https://unstats.un.org/gist/) under the UN. ISLP has also supported GIST in collecting statistical literacy resources in the Statistical Literacy Inita-tives Inventory which can be found at https://unstats.un.org/gist/statistical-literacy.

Periodic meetings are held with the executive team and the advisory board to generate greater interaction in decisions, communication, activities, and representation in different international events. An ISLP Open Meeting is always organised in the World Statistics Congress, and in other conferences. Online meetings for Poster Competition participants have been organised in Latin America since 2021. Moreover, different projects and actions are coordinated regionally by Country Coordinators.

The project promotes and advances statistical literacy by organizing sessions at various scientific conferences, both at International Statistical Institute conferences and elsewhere, as well as within the home country of each participant.

> *"Statistical literacy empowers individuals to make informed decisions in their daily lives, from personal choices to societal issues."*

Challenges and possibilities

When promoting statistical literacy worldwide, the ISLP faces many challenges. Literacy in general - not to mention statistical literacy - is a difficult task. In some less developed countries, there may not be many resources available for statistical education and girls might have unequal opportunities for education. In some other countries, statistical literacy and its importance is not well understood in society. AI brings its own challenges to the promotion of statistical literacy. The ISLP also encounters language barriers. English is just one of the main languages spoken in the world. Sharing best practices across countries and translating teaching materials between languages is an interesting opportunity, but requires funds to do.

ISLP seeks funding continuously. Due to a lack of sustained core funding, the project can currently be coordinated only at a limited level, which makes it more challenging to implement innovations in practice. Adequate funding would enable forward planning, the implementation of concrete initiatives, and stronger support

Figure 2. Source: StatBel

for developing countries. Without stable resourcing, it is not possible to consolidate operations or ensure continuity. Although volunteer contributions have enabled significant achievements, a more secure structural and financial foundation is required to sustain and further develop the project.

A possibility is international collaboration, which offers a powerful and inspiring pathway for strengthening statistical literacy worldwide. Working together with international partners, such as UNITAR/GIST and the ISI Academy, enables access to a wide range of expertise and perspectives.Through such partnerships, it becomes possible to co-develop educational materials, organise joint events, and promote best practices in teaching and communicating statistics.

> *"Adequate funding would enable forward planning, the implementation of concrete initiatives, and stronger support for developing countries."*

To support the ISLP, donations can be made via the ISI donations page: https://isi-web.org/donate

About Reija Helenius

Reija Helenius (b. 1963) has directed the ISI/IASE International Statistical Literacy Project since 2010, during which time a worldwide country coordinator network has been established and the project has grown to operate across all continents. She is Group Manager at Statistics Finland, a board member of the Finnish Statistical Society, and was a member of the ISI Council 2021–2025. She holds a Master's in Social Sciences, a PD diploma in Statistics, and is formally trained in quality management and product development. She has been an ISI Elected Member since 2006, with memberships in IAOS and IASE.

References

Exploristica. Best cooperative project award in statistical literacy 2015: Announcement of winner. [Online]. https://shorturl.at/3GkBS

Falck, E., Campos, P., D'Amelio, A., & Helenius, R. (2025). Global challenges and perspectives in statistical literacy: An analysis of ISLP activities and insights from country coordinators and NSO resources. *Statistical Journal of the IAOS, 41*(3), 940–951.

Helenius, R. (2018, July). International statistical literacy project past and now. In *Looking back, looking forward. Proceedings of the Tenth International Conference on Teaching Statistics (ICOTS10)* (pp. 1–6).

ISLP. *ISLP Competitions* [Online]. https://iase-web.org/islp-competitions

ISLP. *Home* [Online]. https://iase-web.org/islp/

ISLP. *Publications: Newsletters* [Online]. https://iase-web.org/islp-newsletters

MacFeely, S., Campos, P., & Helenius, R. (2017). Key success factors for statistical literacy poster competitions. *Stat Educ Res J, 16*, 202–216.

Statistical Multimedia Literacy Program (LeME). Best cooperative project award in statistical literacy 2019: Announcement of winner. [Online]. https://shorturl.at/hoDtU

The National Schools' Poster Competition. Best cooperative project award in statistical literacy 2017: Announcement of winner. [Online]. https://shorturl.at/fGaci

TREND. Best cooperative project award in statistical literacy 2021: Announcement of winner. [Online]. https://urli.info/1pqlj

What's Going On In This Graph? ISLP competitions: best cooperative project award in statistical literacy 2023: Announcement of winner. [Online]. https://urli.info/1pqlv

Why some countries measure race—and others refuse

Abstract. Official statistics on race and ethnicity are required by law in some countries and explicitly prohibited in others, yet both approaches are defended as advancing social justice and human rights. Jonathan Auerbach explores the history of racial classification, arguing that the debate is not only political but also statistical: official measurement often illuminates injustice, while mismeasurement reinforces it.

Jonathan Auerbach
Jonathan Auerbach is an assistant professor in the Department of Statistics at George Mason University. His research spans the intersection of statistics and public policy, with applications ranging from urban analytics to official statistics, including assessing the quality of the 2020 census and debunking urban myths such as the claim that New York City is home to eight million rats.

Official statistics are the measurements governments take of society. Examples include the population size, the unemployment rate, and life expectancy. These measurements are 'official' because society accepts them as fact for commerce, policymaking, and public discourse.

But why accept these numbers—or any numbers for that matter?

We do not trust official statistics simply on the authority of the government that produced them. Rather, as with any scientific endeavor, we trust the process by which governments balance the competing considerations of data collection. This column examines those considerations.

To collect or not to collect

The first step of data collection is to decide whether to collect data in the first place. Not every aspect of society is suitable for measurement, and countries vary widely in what they choose to measure. Perhaps the starkest example is the decision whether to collect data on race and ethnicity—a fitting topic for this issue on social justice and human rights.

Race and ethnicity are categories used to group individuals according to a shared identity, often based on visible traits such as skin color, cultural heritage, and national origin. In some countries, such as Brazil, South Africa, and the United States, governments are required by law to collect data on race and ethnicity. In others, such as France, Germany, and Japan, collecting such data is severely limited and, in many cases, explicitly prohibited.

Both justify their decision in the name of social justice and human rights. Those in favor of collecting official race and ethnicity data argue that such information is one of the most powerful tools for documenting inequality and discrimination. Those opposed argue that race and ethnicity are at best crude measures of lived experience and at worst categories that reinforce the very inequalities and forms of discrimination their proponents seek to eliminate.

But how does a country decide whether collecting such data advances the public good? To better understand both perspectives, let's consider the origin of race classification with a focus on its use in the United States and its absence in France.

> *"Are countries that refuse to collect race and ethnicity data willfully ignorant? Or are countries that do collect them mistaking an ill-defined classification system for scientific fact?"*

The origin of race as a classification

The history of measuring race is inseparable from the history of racism, the belief that race is a fundamental determinant of human traits or abilities. In this sense, race is a surprisingly modern concept. Ancient civilizations were ethnocentric and xenophobic, and various forms of hierarchy existed. But none appear to have organized around race in the modern sense.

For example, historian George Fredrickson notes that while Roman slavery included people from conquered territories across Africa, Asia, and Europe, enslaved people could in principle become Roman and participate fully in Roman society. Differences of origin did not produce the kind of durable, inherited group identities that warranted an official racial classification system.

Official statistics developed gradually over the centuries of European exploration and colonization, as increasingly centralized states required standardized measurements for the administration of public affairs. Race-like classifications became shorthand for social roles.

For example, colonial administrators in the seventeenth century distinguished among indigenous peoples, predominantly European colonizers, and enslaved populations that were

increasingly of African origin. One reason, as historian James Cassedy explains, was administrative. Colonial authorities sought a "safe" balance, offsetting the increased trafficking of Black slaves with a commensurate amount of White immigration.

As race-like categories became standard, social roles were increasingly interpreted as reflecting fundamental differences between races. The irony was that this view conflicted with the ascendant religious and Enlightenment principles of the period, which emphasized the equality of individuals.

Two revolutions undertaken in the name of those principles reached strikingly different conclusions. France moved, at least in principle, toward a universal citizenship that treated race and ethnicity as irrelevant and later placed strict limits on their measurement. The United States, by contrast, chose to build its constitution around slavery without naming it directly. In doing so, it created a system in which the measurement of race became increasingly necessary for the administrative functions of the state.

The rise and fall of race as a science

The term "race" first referred not to differences among human beings, but to breeds of horses and dogs and to aristocratic bloodlines.

In 1735, Carl Linnaeus placed humans within his taxonomy of plants and animals under the order Primates. Humans were then subdivided into four groups: Americans, Europeans, Asians, and Africans. As biologist Stephen Jay Gould explains, Linnaeus simply organized human beings according to four major geographical divisions, which was not much different than other racial classifications of the time.

Then in 1776, Linnaeus' student Johann Blumenbach revised these categories. Blumenbach not only asserted five groups but also a ranking that placed White "Caucasians" at the top as the ideal from which the other races had "degenerated."

This shift from classification to ranking marked the transition to what is now called scientific racism. The field claimed to learn the laws of biological determinism through measurements ranging from the concrete, such as brain size, to the abstract, such as beauty and intelligence.

Today we know this work was irredeemably flawed, resting on statistical mistakes and, in some cases, outright fraud. Yet for a century and a half it remained enormously influential. Its power lay not only in asserting that race was a principled basis for organizing society, but in co-opting official statistics to make that claim appear objective and representative of the population as a whole.

For example, the 1840 U.S. Census tabulated Americans by both race and insanity, and the results initially appeared to show that free Black Americans in the North were insane at far higher rates than enslaved Black Americans in the South. Historian Patricia Cline Cohen has shown the results were erroneous, likely because the complexity of the census form led enumerators to misclassify senile White Americans as insane Black Americans. But the findings fit the tenets of scientific racism, so the errors were largely dismissed and the conclusions accepted as fact.

The example suggests that science, including the measurement of official statistics, is not automatically self-correcting. Indeed, the collapse of scientific racism owes perhaps less to the triumph of good science over bad than to the horror provoked by the treatment of peoples deemed inferior in the name of science, most notably the mass murder of Jewish Europeans during the Holocaust and the brutal suppression of Black protesters in the American South during the Civil Rights Movement.

It was the moral rejection of biological determinism by ordinary people that did much to clear the way for better science to prevail, rather than the scientific community alone, which in many cases benefited from the preservation of social hierarchy.

But the science did change. In 1972, Richard Lewontin produced what is perhaps the most famous application of the analysis of variance. He showed that humans possess no genes unique to any race and that most genetic variation occurs within, rather than between, racial groups. As statistician A. W. F. Edwards later pointed out, it is not correct to conclude from this work alone that racial classifications have no biological basis. Nevertheless, Lewontin's work helped launch a new body of scholarship demonstrating that the diversity of human traits and abilities cannot be adequately understood through the lens of race and that, in many contexts, race provides little if any biological insight.

Race as a protected identity

Race is no longer accepted as a principled basis for organizing society. Yet racism and discrimination persist. The debate has shifted to the role of government in addressing them. In the United States, official statistics on race have become central to the measurement of discrimination, with mixed success.

For example, in principle, the Civil Rights Act of 1964 prohibits employers from hiring, firing, or setting compensation based on race. In practice, statistician Paul Meier has argued the adversarial structure of the legal process can undermine the objectivity of official statistics as evidence: The plaintiff typically retains an expert witness to analyze employment or wage data and argue that discrimination has occurred, while the defense retains its own expert to analyze the same data using different methods and offer an alternative explanation in which discrimination has not occurred.

Courts are not the only setting in which adversarial relationships politicize ostensibly objective data. In 2005, for example, President George W. Bush removed the head of the federal agency responsible for producing official criminal justice statistics after he refused to downplay the role of race in police stops. Even in relatively apolitical environments, adversarial relationships—from competition in business to academic peer review—shape the interpretation of data. Though less extreme, these contests hearken back to the era of scientific racism, when official statistics, rather than depoliticizing public discourse, were themselves politicized.

"It was the moral rejection of biological determinism by ordinary people that did much to clear the way for better science to prevail, rather than the scientific community alone, which in many cases benefited from the preservation of social hierarchy."

France, by contrast, has largely prohibited the collection of official data on race. Compared with the United States, there is relatively little public discourse about race. But as sociologist Patrick Simon has pointed out, the strategy of "equality through invisibility" has not eliminated racism from French society.

It is difficult to produce a definitive measure of race and racism in France precisely because of the absence of an official statistic. Alternative measures exist, including the use of proxy variables such as names, audit studies, and small-scale surveys. Yet the growth of a political movement calling for the official measurement of race and ethnicity suggests that these tools are seen as insufficient, not necessarily because they are inaccurate, but because they do not provide the commonly accepted basis for public discourse supplied by official statistics.

The limits of measurement

Are countries that refuse to collect race and ethnicity data willfully ignorant? Or are countries that do collect them mistaking an ill-defined classification system for scientific fact?

One's perspective depends in part on values, such as the conflicting ideals of American multiculturalism and French universalism. But history suggests that the issue is not only political. It raises genuine statistical questions about whether a concept can be measured meaningfully, and whether it ought to be measured at all.

As the statistician William Kruskal argued, race and ethnicity are attempts to measure real if abstract concepts, and therefore statisticians have an important role to play in evaluating them. Yet this does not mean that everything measurable should be measured. As another statistician W. Edwards Deming warned, "It is wrong to suppose that if you can't measure it, you can't manage it—a costly myth."

References

Cassedy, J. H. (1969). *Demography in Early America: Beginnings of the Statistical Mind, 1600–1800.* Harvard University Press.

Cohen, P. C. (2016). *A calculating people: The spread of numeracy in early America.* Routledge.

Deming, W. E. (2018). *The new economics for industry, government, education.* MIT Press.

Edwards, A. W. (2003). Human genetic diversity: Lewontin's fallacy. *BioEssays*, 25(8):798–801.

Fredrickson, G. M. (2002). *Racism: A Short History.* Princeton University Press.

Gould, S. J. (1981). *The Mismeasure of Man.* W. W. Norton.

Kruskal, W. (1981). Statistics in society: Problems unsolved and unformulated. *Journal of the American Statistical Association*, 76(375):505–515.

Lewontin, R. C. (1972). The apportionment of human diversity. *Evolutionary Biology*, 6:381–398.

Meier, P. (1986). Damned liars and expert witnesses. *Journal of the American Statistical Association*, 81(394):269–276.

Simon, P. (2008). The choice of ignorance: The detabe of ethnic and racial statistics in france. *French Politics, Culture & Society*, 26(1):7–31.

> *"Indeed, the collapse of scientific racism owes perhaps less to the triumph of good science over bad than to the horror provoked by the treatment of peoples deemed inferior in the name of science, most notably the mass murder of Jewish Europeans during the Holocaust and the brutal suppression of Black protesters in the American South during the Civil Rights Movement. "*

An Interview with Tony Labillois

On official statistics, accessibility, and the responsibility of data

Tony Labillois is a Canadian statistician with more than 35 years of experience in official statistics, accessibility advocacy, and international cooperation. Rising from methodologist to Director General at Statistics Canada, he co-led the Disaggregated Data Action Plan, served for over two decades as Champion for Persons with Disabilities, and represented Canada before the United Nations Human Rights Council. In 2025, he was recognised as an Elected Member of the International Statistical Institute. In this interview, we trace his remarkable journey, from a small village in the Gaspésie to the halls of Geneva, and explore what it truly means for statistics to be inclusive, accessible, and impactful.

Role: Director General, Statistics Canada (ret. 2024) **ISI:** Elected Member 2025 **Advocacy:** Champion for Persons with Disabilities 2002–2024 **Honours:** Diamond Jubilee Medal (2012); Coronation Medal (2025)

Manuele: *First of all, congratulations on becoming an elected member of the ISI. How did you feel when you received the news?*

Tony: I was very honored, almost speechless. It made me reflect on my career and my achievements, and on the fact that other esteemed colleagues from different countries and different vantage points would support me as an elected member of such a prestigious and historic organization. It meant a great deal.

At the same time, I was deeply humbled, because it made me think of so many people who were there along the way, giving me advice, encouraging me to develop my abilities, celebrating the steps with me. Whatever I have achieved belongs as much to those encounters and collaborations as it does to me.

Manuele: *Is there anyone in particular who really shaped your journey?*

Tony: My first director, Dr. Boverianda Nanjamma Chinnappa, shaped my leadership significantly. Then people like Eric Rancourt, someone I met in my early years at Statistics Canada and still a friend today. Pierre Lavallée, another elected member, was one of the first to encourage me to present at a conference, just one year after I joined. And then Yves Tillé, another supporter, and Yves, whom I have known since he came to Canada for a colloquium I was organizing in Rimouski. And later, Steve MacFeely from the OECD.

This honor gave me a great deal of energy and a strong desire to give back. I have received so much in my life, and I want to keep serving, not only with statisticians, but also with people who do not yet appreciate the importance of data and statistics in our digital world. Recognition is meaningful only if you leverage it to do more, as others did for me. Now that I have more time since my retirement from Statistics Canada in July 2024, the timing is good to pursue things that will hopefully have a positive impact, alongside people who share my values and a form of servant leadership.

Manuele: *To share a little of your story that others may not know: you have mentioned that you were born with low vision. Perhaps there is more understanding and inclusivity now, but looking back, what was your experience as a student of mathematics and statistics?*

Tony: I was born in a small village between the mountains and the sea in eastern Quebec, the Gaspésie region near New Brunswick, in the late 1960s. People were asking my parents, "What will you do with Tony? You should send him to an institution in Montreal," which was over a thousand kilometers away, more than ten hours at the time. But my parents believed that everything was possible and that I could find my own path. They wanted to send me to the regular school in the village with the other children, and eventually they got agreement from the principal and teachers.

I grew up there, went through high school, and realized very quickly that I was not only good at mathematics and science but genuinely enjoyed them. There were no accommodations at the time, except perhaps sitting at the front of the classroom to see the blackboard, or asking a friend what was written there. Those were also the days of stencil duplicators, and they would give me the darker copies when possible. My nose would end up slightly blue from the ink because I held the paper so close to read it.

> *"He made me cry. But I told him I still loved it and was still confident I would eventually find a fruitful path."*

When I arrived in Quebec City for college, I had two years of study that included twelve science classes, six of which were mathematics, and my first statistics course. I also had my first personal computer in 1981 and taught myself to program. Given my low vision, I had realized that I wanted a job I could perform with full satisfaction and autonomy. Chemistry or microbiology would have meant difficult lab instruments; engineering, construction sites or factories. Statistics and computing felt like a fruitful path.

So I went to university and specialized in statistics. That was hard. I had a small telescope to see graphs, but professors would not share their notes, there were no accommodations for exams, and at some point you simply need to see the formulae being drawn on the board. Then one professor took me aside and told me I had no place there, that there were too many things I could not see, and that I should consider something else entirely. He made me cry. But I told him I still loved it and was still confident I would eventually find a fruitful path.

Despite that discouragement, I was successful in my classes. It was only at Statistics Canada that I eventually found my proper place. They were far more welcoming than university had been.

Manuele: *This story is both sad because of the difficulties, and remarkable because of the way you managed to push through them and build such a successful career. And perhaps those very difficulties are what motivated you to become such a strong advocate for inclusivity.*

Tony: Perhaps. I will say that I was born with low vision, so for me, this is my normal eyesight, the same I have had my whole life. I can still travel alone, make photographs, enjoy art, and participate fully in the world of sight. When I look at data, I need a graph. I understand things much better through charts and images, and I draw on a board when working in a team.

> "Accessible data is not just a compliance exercise. It is a design choice and a discipline in itself."

What I only realized when I finished secondary school was that perhaps I should not have been so ashamed of myself. The school I attended instituted a prize in my name, and I was its first recipient. That changed the way I saw myself. During that summer between high school and college, I became more confident, more inclined to explore my abilities, and I realized I had already developed some. That gave me strength, and it is another example of why I want to give back.

When I arrived at Statistics Canada, my first chief and first director offered me accommodations, even before I was hired. The application form asked candidates what accommodations, if any, they would need, which was novel in 1989. I mentioned large print for the exam and a bit of extra time given the volume of reading involved.

I arrived at the exam room on a Saturday morning to find my accommodations were not ready. I said, "Give it to me anyway, I am used to managing without." But the person there insisted I wait, and contacted me a few weeks later to sit the exam in a room with better lighting, large print materials, and a little extra time. To my genuine surprise, I passed. They invited me to an interview, again with large print questions, and when I was hired, they told me: "You are hired for your competencies and your potential. We want to make sure you have all the tools you need, just as anyone else would, to achieve results for yourself and to contribute to the team."

For me, that was a revelation. I felt accepted, included, welcomed, and respected. That was another source of motivation to eventually give back, though for many years I simply got on with the work. I was happily surprised by the trust they placed in me, to supervise teams, lead interesting projects, and eventually to be promoted. I started as a recruit and finished as Director General, with many people along the way encouraging me to keep learning and explore further opportunities. I did not really begin advocating actively for accessibility until 2002.

Manuele: *You have described a remarkably varied career at Statistics Canada. Across all of it, what did you find most exciting or meaningful?*

Tony: The beauty of an organization like Statistics Canada is the opportunity to develop yourself, move around, and work with knowledgeable people across every domain: economic, social, demographic, environmental, macroeconomic, and not just in statistics and data, but across the policies, programs, and services connected to economy, environment, and society. When I became a chief in 1997, I led special business surveys from start to finish, surveys that did not exist in the base program, commissioned by clients willing to pay for research they needed: aerospace, energy consumption in commercial buildings, the Y2K bug. That kind of continuous learning went on for decades.

In official statistics, you get to be impactful. You contribute to data and insights that inform better decisions, policies, and legislation for the public good, economic prosperity, social cohesion, quality of life, and human rights.

One example: when I was responsible for the softwood lumber statistics project, during a dispute between Canada and the United States over tariffs, I got to produce statistics at an incredible pace, drawing on a team from across Statistics Canada, to answer the needs of lawyers representing Canada, working with provincial and federal governments and dozens of stakeholders. For me, coming from a village whose economy was built around forestry and sawmills, I could suddenly connect my roots directly to my work.

Another example: as Director General, I co-led the Disaggregated Data Action Plan, collecting more detailed information about different segments of Canada's population. I could see us changing the culture at Statistics Canada so that instead of just looking at national averages, analysts would look at what was happening for racialized groups, women, Indigenous people, and persons with disabilities, across gender, crime, immigration, employment, income, and so on. I even got to represent Canada before the Human Rights Commission in Geneva and speak about how statistics can be a mirror for democracy.

What I find most meaningful is seeing how official statistics can make a real difference through partnerships,

Tony Labillois at the Statistics Canada presentation of new census data on Canada's Hispanic-Latin American population. Photo: Isabel Inclan.

lasting programs, and above all lasting relationships. Data is a lever for change, and official statistics are becoming ever more foundational in the digital era, for digital sovereignty and responsible, grounded AI.

You can probably hear my passion as I speak about this.

Manuele: *Official statistics is sometimes undervalued in academia. If you had to convince a recent graduate to join a national statistics agency, what would you tell them?*

Tony: I would allude to the examples I have just shared, and add this: official statistics matter because they help us understand the world as it really is, not just as we think it is, or as we believe it should be. They allow us to break out of our echo chambers and understand what is truly happening, through neutral, rigorous, independent, and transparent quantitative information, complemented by qualitative information, of course, but the quantitative work is absolutely necessary.

The coherence and consistency, geographically, over time, and across disaggregated segments of our economy, population, and environment, is extremely important. These statistics provide trusted, objective information that helps governments, businesses, and communities make much better decisions about jobs, health, education, the environment, things every citizen should care about regardless of their background or comfort with numbers.

> *"Official statistics turn data into understanding, and understanding into action for the public good."*

Without official statistics, we would be relying on guesses, opinions, and incomplete information. With them, we have a shared foundation of facts that supports fairness, accountability, and better outcomes for everyone. And when you do not measure something, it tends not to exist in people's minds. But when you measure well and make the results known and understandable, you have a very strong lever for moving your society, your economy, your environment toward a better future.

Official statistics turn data into understanding, and understanding into action for the public good.

We need the brightest minds contributing to the tangible problems we face at this pivotal moment. With the advent of AI, all the challenges we faced in the past are exacerbated. AI is a catalyst: it reveals biases, inconsistencies, lack of interoperability, lack of data governance, algorithms that need to be more transparent. We are at an opportunistic moment to influence what happens next. The status quo is no longer acceptable.

Manuele: *Statistics Canada clearly had accessibility in mind from your very hiring process. What do you think needs to happen to make the data world more broadly inclusive and accessible, and what does it actually mean for data to be accessible?*

Tony: There are actually two ways to take that question, and they connect in the end. You can think of data that is accessible by everyone, or data on accessibility itself.

On the first dimension: data that is truly accessible is data that people can find, understand, and use in their day-to-day work, regardless of their ability or background. A lawyer is not used to reading data tables the way we are. My mother certainly is not, unless we digest the data into a clear form of information highlights. So we have a role to play in making data available in clear language, well-documented with metadata, in interoperable formats, and on platforms that meet accessibility standards.

I spoke about this just yesterday in front of 1,200 civil servants at the Canada School of Public Service, about designing for different users from the very start, and thinking about data from the beginning of anything we design, whether that is a policy, a program, or a process. The data needs to be not just available, but truly usable and used.

Going back to the softwood lumber example: the lawyers did not understand the data tables we produced. When we realized that, we restructured the output to follow the same flow as the questionnaire, logging, lumber, wood products, and inserted derived variables in between. Suddenly

Tony Labillois speaking at the United Nations Human Rights Council during the Fourth Universal Periodic Review of Canada.

they understood exactly what they had in their hands and could explain it to other lawyers and stakeholders. That is when it becomes a lever. It took effort for us to realize they did not understand, and then they flew us to Washington to explain the tool to others. Accessibility of data in practice is extremely powerful when it all comes together.

On the second dimension, data on accessibility itself: you need data that makes visible what is often invisible, that captures the lived experience of persons with disabilities, that measures barriers and not just outcomes. What makes someone like me able to achieve results is how I was able to overcome barriers, and how my employer and society more broadly helped remove them. That requires thoughtful, respectful data collection, disaggregation, and ongoing engagement with communities to ensure the data reflects reality.

Accessible data is not just a compliance exercise. It is a design choice and a discipline in itself. It means asking, every time: What is this data for? Who might be excluded? What would it take to include them? If you can answer those questions consistently, you move from data that merely exists to data that makes a real difference, in discussions, decisions, and actions for public good, prosperity, quality of life, and human rights.

This also helps make data more valued by people who do not traditionally see the importance of statistics, and that is increasingly urgent now that people are using AI without realizing it can generate hallucinations, sending them in the wrong direction in their decisions and actions.

Manuele: *You have shared that you are now retired, and as you move into this new phase of your life, will you finally slow down, or is that unlikely?*

Tony: That really makes me laugh, because my wife would tell you I will probably never slow down.

I am still filled with energy, passion, curiosity, a sense of joy and awe, and a very strong desire to serve, to give back, to transmit the values I have received, and to keep learning. When I have more time, I want to maximize the impact. I hope to stay healthy, relevant, and impactful, and humbly, to be both inspiring and inspired by others.

I think of my father. He passed away three years ago. In those last days we spent with him, I could see how proud he was, how satisfied, how at peace with his accomplishments and his relationships. We all know how this ends, and I think there is really only one good thing to do: walk the path together, because we are stronger together, and maximize the positive impact we can have from the values we receive and transmit.

The ISI is a great organization for exactly this: a network, a professional community, or really several communities. As we work together, we can influence the future of our nations for the better. I have worked in the trenches for many years, but when you lift your gaze a little further ahead, you can bring something to the long-term vision, something that keeps our policymakers, leaders, and citizens focused on reality, on challenges, and on solutions for the greater good.

That is why I am so excited. It is much more interesting than finishing the next paper, or doing something pseudo-urgent that will be superseded in a few days.

Manuele: *Thank you so much, Tony. This was truly inspiring. And thank you for your service to the statistics community, and for sharing your story with us in vISIon.*

Tony: This is my pleasure. It is an honor and a privilege. Looking back at my little story, from the small boy from the Gaspésie who was uncertain about his future…I really am an outlier. And instead of being uncomfortable with that, I want to celebrate it. We should celebrate diversity in every possible sense, especially when it leads not just to personal achievement, but to the collective success of our societies, our economies, and our lives. For me, that is what connects human rights, a sense of belonging, and individual achievement, because we are all very different, yet so similar in our aspirations.

Building computational workflows for human rights archives

Design decisions under data sensitivity constraints

Abstract. Computing tools for managing and analyzing historical archives related to human rights face a design paradox that is rarely discussed in the literature on software: the data that drives their development cannot be made public. This article describes a two-layer architecture that has emerged from ongoing academic projects in Uruguay and Chile, in which natural language processing is applied to declassified military and government archives from the Southern Cone dictatorships of the 1970s. The first layer includes tools for information retrieval and document structuring—optical character recognition (OCR) on poorly preserved scanned documents, named entity recognition, and information extraction processes—whose logic can be described and partially demonstrated, even when the underlying corpus is restricted. The second layer comprises search and aggregate analysis tools that allow transitional justice researchers and the general public to query large-scale, structured archival data. We analyze the methodological decisions imposed by sensitivity constraints, the evaluation challenges they pose, and what these projects suggest for the growing field of computational approaches to human rights documentation.

Elina Gómez
Assistant Professor at the Instituto de Investigación en Justicia Social y Desigualdades, Universidad de la República (Udelar). She is a PhD candidate in Sociology at the Facultad de Ciencias Sociales, Udelar. Her research focuses on computational social science, large-scale text analysis, and statistical learning applied to social research and public policy design.

Lorena Etcheverry
Associate Professor at the Instituto de Computación, Facultad de Ingeniería, Universidad de la República (UdeLaR), where she coordinates the Master's Program in Data Science and Machine Learning and leads the Responsible Data Science area at CICADA.uy. Her research covers knowledge graphs, data quality, privacy-preserving analytics, and the detection and mitigation of algorithmic bias.

Jocelyn Dunstan
Associate Professor with a joint appointment in Computer Science and the Institute of Mathematical and Computational Engineering at Pontificia Universidad Católica de Chile. She holds a Ph.D. in Applied Mathematics from the University of Cambridge. Her research focuses on natural language processing and machine learning applied to clinical narratives and public health surveillance.

Introduction

Uruguay and Chile experienced military dictatorships of exceptional brutality and duration. In Uruguay, the civic-military dictatorship lasted from 1973 to 1985; in Chile, the Pinochet regime came to power on September 11, 1973, and lasted until 1990. Both regimes systematically violated human rights through detention, torture, forced disappearance, and exile, leaving tens of thousands of victims. Decades later, these periods remain the subject of active historical debate, ongoing judicial proceedings, and social mobilization in both societies.

A vast documentary record of this period exists in archives that remain partially inaccessible, scattered among institutions in different countries, or restricted by legal and administrative barriers. These archives are of critical importance both for national historical research and for individual and family histories that remain unresolved to this day — families still searching for missing relatives, survivors seeking recognition, and courts still processing cases. In the case of Chile, a modification the law No. 18.771 in January 1989 allowed for the elimination of ministerial documents without authorization

from the National Archives, leading to the estimated destruction of massive amounts of institutional and military documentation just months before the regime ended.

But the software built for human rights archives in the Southern Cone faces a striking asymmetry. Making archival data public brings enormous collective benefits for research, justice, and memory. However, this openness must be weighed against the urgent need to protect individuals—victims, witnesses, and even perpetrators—who often have their personal information included in these sensitive documents. Archives from Uruguay and Chile's 1970s dictatorships contain details that, if freely disclosed, could endanger lives or disrupt ongoing judicial processes Lessa (2022); Etcheverry et al. (2021). Still, constructing effective computational tools for transitional justice depends on using precisely these documents. Here, the archive is double-edged: at every step, decisions on access, storage, and inference become ethical choices.

This article is itself a product of that tension. We describe two related projects — the Memorias project at the Universidad de la República (Uruguay) and the NuestraMemorIA initiative, a joint effort between the Pontificia Universidad Católica de Chile and the Universidad Alberto Hurtado (Chile) Díaz et al. (2024) — and the common design space they share. Our primary goal is to clarify the architectural and methodological choices required when designing natural language processing pipelines for archives containing sensitive data. Rather than benchmarking models, our aim is to provide a focused account of the challenges and decisions that shape the development and operation of computational tools in this unique context.

Data Layers

The tools we have developed in both projects are naturally organized into two layers that serve different audiences under different constraints.

Layer 1: Document Recovery and Structuring

The first layer addresses the raw digitization challenge. The Southern Cone repressive agencies produced documents that were typewritten on poor-quality paper, mimeographed, annotated by hand, and preserved under adverse conditions for decades. The performance of standard off-the-shelf OCR systems — even current vision-language models — significantly degrades on this material Rang et al. (2024); Sastre et al. (2025).

The Uruguayan *Archivo Berrutti*, for example, consists of millions of pages from the Dirección Nacional de Información e Inteligencia (DNII), of which tens of thousands have been digitized and collaboratively annotated to date Etcheverry et al. (2021). State-of-the-art OCR systems achieve around 98% character accuracy on standard printed Latin-script documents Rang et al. (2024), but performance degrades significantly on degraded historical material, with errors concentrating on proper names, dates, and handwritten annotations Sastre et al. (2025).

The general flow proceeds in three stages:

1. **Preprocessing.** Scanning artifacts are removed and page layout is analyzed to detect regions of interest (text blocks, stamps, handwritten fields).

2. **Recognition.** A locally deployed, fine-tuned model produces character-level transcriptions.

3. **Post-processing.** Domain-specific rules normalize dates, reference numbers, and named entities before storage in the structured repository.

A key design decision here is local inference. Both projects process documents on the institutional infrastructure rather than commercial cloud APIs. This is not merely a cost decision: it is a legal and ethical requirement. Sending archival documents containing personal data about victims or alleged perpetrators to external servers would conflict with data protection regulations in both countries and, more importantly, with the trust placed in us by families and human rights organizations.

Named entity recognition (NER) follows OCR and presents its own challenges. The entities we care about — people, military units, geographic locations, dates — appear in ambiguous, context-dependent forms characteristic of intelligence documents. The Uruguayan team has approached this through a combination of fine-tuned transformer models for named entity recognition Gallardo Negrín (2025) and knowledge graph construction with human-in-the-loop annotation to manage uncertainty Díaz et al. (2024). A prerequisite for this entire layer is archival expertise. The typology of documents — distinguishing a *parte diario* from an intelligence *ficha*, a surveillance report from an interrogation record — cannot be inferred automatically and must be established upstream. Archivist knowledge shapes the ontology, which in turn shapes the NER schema and the knowledge graph structure.

> *"The records that most demand computational attention are precisely the ones that most resist standard computational workflows."*

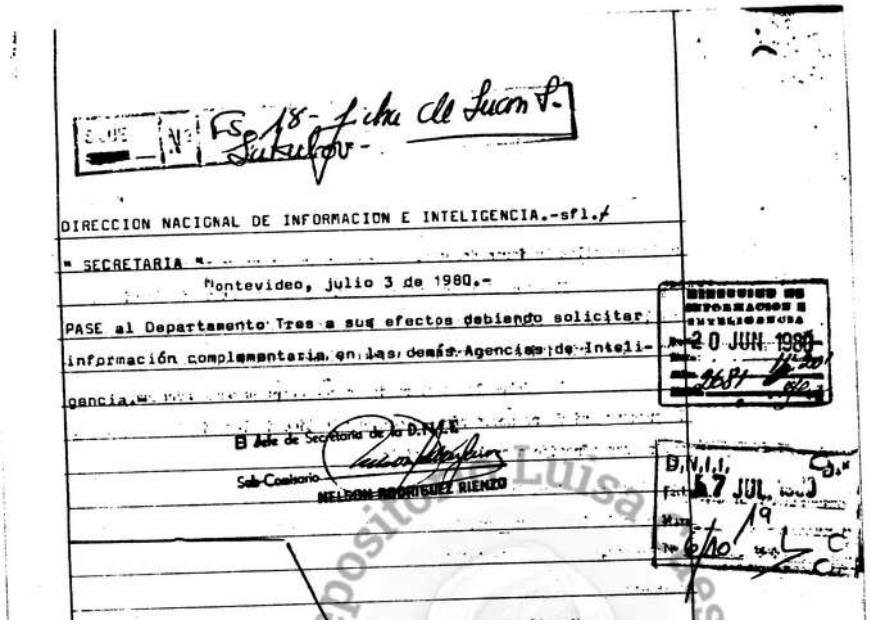

Figure 1. An original intelligence document from the Archivo Berrutti (Dirección Nacional de Información e Inteligencia, Montevideo, 1980). Overlapping stamps, handwritten annotations, and degraded typeface are characteristic of the OCR and layout challenges described in this section.

In the Chilean experience of document recovery and structuring, researchers have addressed the failure of standard OCR systems to handle the severe physical degradation, faded ink, and domain-specific lexicon characteristic of dictatorship-era archives. To overcome these challenges while ensuring data sovereignty and privacy, the team developed HiT (Here is the Text), a locally deployable hybrid strategy that avoids reliance on external cloud services. This two-stage method first involves fine-tuning specialized Spanish text recognizers on newly created, large-scale datasets—DHiSS and DHiSS+—which mirror the specific typographic conventions and noise patterns of the target material. In the second stage, high-confidence lexical *anchors* from these recognizers are injected into a Vision-Language Model (VLM) to ground its text generation, reducing in this way hallucinations Vasquez et al. (2026).

Layer 2: Search and Aggregated Analysis

The second layer transforms structured archival data into queryable interfaces. Here, the audiences diverge: academic researchers in history, political science, and transitional justice need fine-grained document-level queries; the broader public and human rights organizations need to ask questions like "what documents mention this person?" or "which military units appear in connection with this region in 1975?". In the Uruguayan project, this layer currently comprises two tools: LUZ, a search tool based on a relational database that allows users to perform full-text searches on document transcripts, and Amalia, a platform for qualitative content analysis and researcher-assisted interpretation of archival material Etcheverry et al. (2021); Gómez Bonaglia (2026).

In the Chilean project, this layer is operationalized through GraphRAG, an automated framework that couples retrieval-augmented generation (RAG) with knowledge graph construction to transform noisy corpora into interactive network representations . This system enables users to interact with fragmented archives via a natural language chatbot capable of answering complex queries grounded in explicit graph evidence . By integrating Cypher property queries and hybrid search, the platform effectively reconciles conflicting testimonies into coherent narratives, as demonstrated in case studies involving the Orlando Letelier assassination and the disappearance of Juan Luis Quiñones Sanhueza et al. (2026)

Both projects share a commitment to **sovereign retrieval**: query systems that run on local models, without sending user queries or document contents to external services. This shapes every infrastructure decision, from model selection to the hardware required.

The Evaluation Problem

Standard natural language processing (NLP) benchmarking assumes the existence of annotated test sets that can be shared alongside model weights. For our domain, this assumption fails at multiple levels. We cannot release raw documents; we cannot always release annotations either, because the structure of an annotated corpus can reveal which documents were processed

and what entities they contain.

We have addressed this through three strategies. First, we construct **synthetic evaluation corpora**: documents generated with the statistical properties of the real archive (similar vocabulary, name distributions, document types) but containing no actual personal information. Second, for the OCR layer, we use **public subsets** — portions of each archive already released under open-access protocols — to report performance numbers, acknowledging that these are upper bounds on what we see in the full restricted corpus. Third, we rely on **expert-in-the-loop validation**: domain historians and archivists review system outputs as part of their research workflows, and their corrections inform model improvement.

This last strategy, we argue, is not merely a workaround for the evaluation problem but a substantive feature. Historical documents require contextual interpretation that automated metrics cannot capture. A NER system that identifies a person's name but misattributes their role in the repressive apparatus is more dangerous than one that simply fails to extract the entity. Quality in this domain means calibrated uncertainty as much as raw recall Jurafsky and Martin (2008).

Results and Shared Lessons

The following example illustrates how the two-layer architecture operates in practice before we draw the general lessons that apply across both projects.

A Concrete Example: Reconstructing a Detention Trajectory

Consider a researcher investigating possible associations between political detention and cancer among former prisoners of Uruguay's dictatorship — a line of inquiry that arises directly from survivor testimony Jorquera et al. (2020). She begins with a keyword search in LUZ using the term *tumor* across the transcribed collections of the Archivo Berrutti. The search returns 47 document images containing that string, but the raw results are of limited use: some are false positives caused by OCR errors *tumor* misread from *turno* or *humor*), and the documents themselves are unstructured images of handwritten or typed pages.

The pipeline then applies two complementary tools. First, domain historians identify documents containing tabular structures — detention lists of the kind shown in Figure 2 — that record names, identity numbers, detention dates, and political affiliations. This classification requires archival expertise and cannot be reliably automated; it is precisely the kind of upstream human judgment described in principle (b) above. Once identified, automated table extraction recovers the structured data locked inside degraded scanned images. Second, a named entity recognition model identifies person names appearing both in the keyword-matched documents and in the detention tables, enabling record linkage across sources: a name mentioned in a handwritten medical note can be tentatively connected to a row in a detention list, and from there to a date of arrest and a political affiliation.

The result is a partial life trajectory — not a complete biography, but a documented sequence: *detained on this date, held in this facility, mentioned in a medical context on this later date.* For the researcher, this is not a conclusion but a hypothesis to be verified against other sources. From a system perspective, this demonstrates that computational tools can uncover connections between diverse and damaged documents that no manual, page-by-page reading could detect efficiently at scale.

This example is illustrative, but the underlying pipeline is not. The detention table in Figure 2 is a real document from the Berrutti/OCOA collections. The same workflow applies to any research question expressible as a search term — a person's name, an institutional acronym, a date range, a medical term — producing structured, linkable data from sources that were previously opaque to systematic analysis.

Across both projects, several design principles have proven robust.

a. **Local-first, open-source.** Every component that can run on institutional hardware, does.

b. **Archival expertise as upstream constraint.** As described in Section 2.1, document typology is a prerequisite for system design, not a downstream annotation task.

c. **Sensitivity as a spectrum, not a binary.** Not all documents in these archives are equally sensitive. Some are already public; others contain information about living individuals whose safety is at risk. Access control must be document-level, not collection-level, and the infrastructure must support graduated access: open search over metadata, controlled access to transcriptions, restricted access to raw images Hernández-Muñiz and López-Carrato (2022).

d. **Evaluation by use.** The most meaningful measure of system quality is whether historians and transitional justice researchers can answer questions with it that they could not answer before. We have adopted a "research-as-evaluation" approach in which each new platform version is validated by real investigations rather than held-out test sets.

> *"At every step, decisions on access, storage, and inference become ethical choices."*

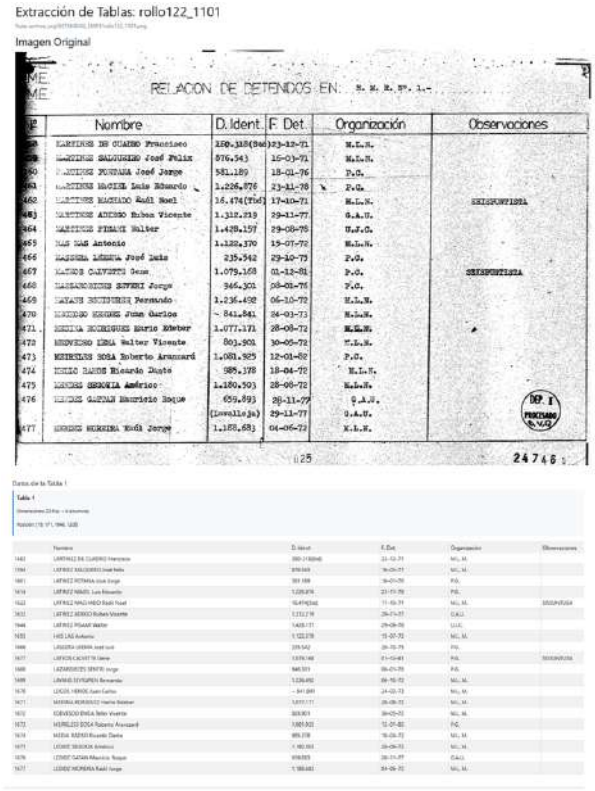

Figure 2. A detention list from the Berrutti/OCOA archives (top) and its automatic structured extraction (bottom). Columns record name, identity number, detention date, political organization, and observations. Table extraction of this kind is the entry point for the record-linkage pipeline described in this section.

Conclusion

The computational tools we have described here are not remarkable for their technical novelty. OCR, NER, knowledge graphs, and retrieval-augmented generation are all established methods. What is novel is the constellation of constraints under which they must be deployed, and the design decisions those constraints force.

These experiences also demonstrate that analyzing sensitive human rights archives requires a transdisciplinary approach that balances computational scalability with ethical rigor and historical accuracy. By co-designing with key institutions, we can move towards the support of the five pillars of transitional justice—truth, justice, reparations, memory, and non-repetition—by turning fragmented, sensitive records into accessible tools for accountability and collective memory.

We hope this report is useful to the growing community working on computational approaches to human rights documentation across Latin America, Southeast Asia, and beyond. The Southern Cone archives are a specific case, but the data sensitivity paradox is general: the records that most demand computational attention are precisely the ones that most resist standard computational workflows.

References

Díaz, C. et al. (2024). Automatic knowledge-graph creation from historical documents: The Chilean dictatorship as a case study. *arXiv preprint arXiv:2408.11975*.

Etcheverry, L. et al. (2021). A computational framework for the analysis of the uruguayan dictatorship archives. In *Qurator 2021 — Conference on Digital Curation Technologies*, Berlin, Germany.

Gallardo Negrín, R. (2025). Extracción de información utilizando modelos generativos en documentos del pasado reciente. Master's thesis, Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay.

Gómez Bonaglia, E. (2026). Las huellas de la represión. Recuperación y procesamiento de archivos militares producidos durante la dictadura uruguaya. In Afanador-Llach, M. J. and Vaughan, N., editors, *Ctrl + F.* Ediciones Uniandes, Bogotá.

Hernández-Muñiz, F. and López-Carrato, M. d. L. (2022). Función social en los archivos de derechos humanos a través del Proyecto Cruzar.uy. *Ibersid*, 16(1):65–74.

Jorquera, M. J. et al. (2020). Mortality in survivors of torture and political imprisonment in Chile. *Revista Médica de Chile*, 148(12):1773–1780.

Jurafsky, D. and Martin, J. H. (2008). *Speech and Language Processing*. Prentice Hall, 2nd edition.

Lessa, F. (2022). *The Condor Trials: Transnational Repression and Human Rights in South America*. Yale University Press.

Rang, M., Bi, Z., Liu, C., Wang, Y., and Han, K. (2024). An empirical study of scaling law for OCR. *arXiv preprint arXiv:2401.00028*.

Sanhueza, D., Mery, D., Fonck, A., Rojas, H., and Dunstan, J. (2026). GraphRAG for the chilean dictatorship: A knowledge graph approach to human rights archives. Under submission.

Sastre, I. et al. (2025). Post-OCR correction using large language models with constrained decoding. *preprint*.

Vasquez, C., Mery, D., and Dunstan, J. (2026). HiT: A hybrid OCR method for transcription of historical documents. Under submission.

Beneath the surface

What data reveal about climate pressure on Indonesian fisheries

Abstract. Indonesia's fisheries are often described through production, exports, and the number of people working at sea. Yet beneath these figures lies a more fragile story: climate change is reshaping the ocean, and with it the livelihoods of coastal workers and fishing communities. This article shows how integrated data by combining official statistics, vessel activity, and environmental or geospatial indicators, can help reveal the hidden links between climate signals and socio-economic vulnerability. The Indonesian case illustrates why future official statistics need to move beyond separate datasets toward evidence systems that connect environmental change with people's lives. In a warming world, measuring climate impact is not only about tracking physical change. It is also about identifying who is most at risk, where support is needed, and how data can help turn visibility into protection.

Suryo Adi Rakhmawan
Senior statistician at BPS–Statistics Indonesia, working at the intersection of population, labour, migration, and official statistics modernization. His recent work explores how official statistics can be strengthened through geospatial data, CRVS, remote sensing, and alternative data sources to support more inclusive and evidence-based policy.

I Nyoman Setiawan
Computer specialist at BPS–Statistics Indonesia whose work focuses on statistical modelling, computational statistics, and applied policy analysis. His recent work explores how geospatial and alternative data sources can complement official statistics in understanding climate-related risks, fisheries, and coastal livelihoods in Indonesia.

Indonesia lives with the sea

For Indonesia, the sea is not a distant landscape. It is a workplace, a food system, a transport corridor, and a source of identity for millions of people. Across the archipelago, coastal communities do not simply live near the ocean; they live with it. But as the climate changes, the sea is becoming less predictable.

Rising sea temperatures, shifting rainfall patterns, stronger weather variability, and changing marine ecosystems are no longer abstract environmental concerns. For many coastal workers, these changes may be experienced as fewer safe days at sea, longer trips to fishing grounds, higher operational costs, unstable catches, and uncertain household income.

The challenge is that these impacts are often hidden beneath conventional statistics. Fisheries may still appear in national accounts, employment tables, or production reports. Yet these numbers do not always show how climate pressure moves through people's lives. The question is therefore **who becomes vulnerable when the sea changes?**

The hidden problem beneath fisheries statistics

Fisheries are commonly measured through visible indicators: production volume, export value, number of vessels, number of workers, or contribution to economic output. These indicators are important, but they cannot fully capture livelihood vulnerability.

Climate pressure often appears gradually. It may not immediately look like a sectoral collapse. Instead, it may emerge as a chain of disruptions: fishing grounds become less reliable, trips become longer, fuel costs rise, catches fluctuate, income becomes unstable, and households adjust their consumption or coping strategies. In this sense, the impact of climate change on fisheries is rarely a single number. It is a pathway from environmental change to human vulnerability.

A warmer sea may affect fish distribution. Changed fish distribution may alter fishing activity. Altered fishing activity may affect working hours, income, and household welfare. In communities where work is informal and social protection is limited, these pressures can become deeper forms of vulnerability.

Why one dataset is not enough

Understanding this pathway requires more than one source of data (Table 1). Official statistics remain essential. Labour force surveys can show who works in fisheries, whether they are formal or informal, and how employment conditions differ across regions. Household surveys can reveal welfare, poverty, consumption, food security, and access to basic services. But official statistics alone may not fully capture what is happening at sea.

Vessel activity data can provide signals of fishing movement and operational intensity. Remote sensing can help identify environmental pressure, such as sea surface temperature, rainfall, wind, or coastal exposure. Geospatial data can show where communities are located, how accessible they are, and whether they are close to services, ports, markets, or hazard-prone areas.

In this study, three environmental variables were mapped across In-

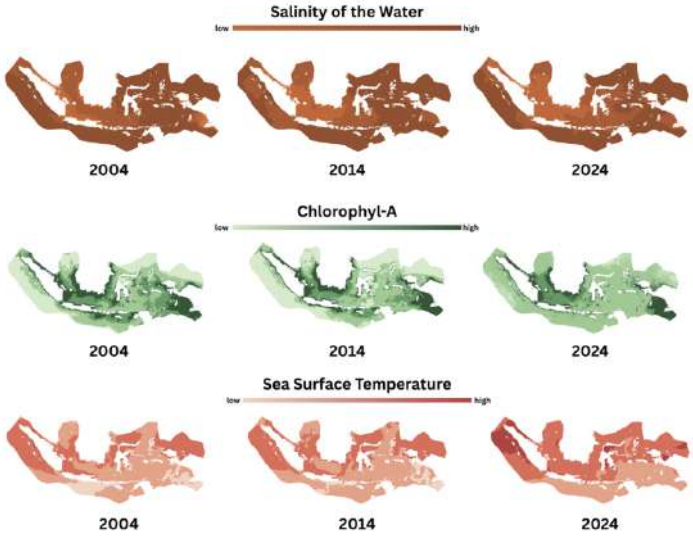

Figure 1. Spatial distribution of environmental variables in Indonesian waters (2004–2024): (a) Salinity of the water (PSU), (b) Chlorophyll-A (mg/m^3), (c) SST (°C)

donesian waters for 2004, 2014, and 2024: sea surface temperature, salinity, and chlorophyll-a (Figure 1). Together, these indicators help capture ocean warming, changes in water characteristics, and marine productivity. Their spatial distribution shows that climate pressure is not uniform across Indonesian waters, making geospatial evidence essential for understanding where fisheries and coastal livelihoods may be most at risk.

No single dataset can tell the full story. But together, they can connect climate signals at sea with livelihood risks on land.

From climate signal to livelihood risk

The relationship can be understood as the simple pathway shown in Figure 2. This pathway matters because it connects macro and micro perspectives. At the macro level, climate change affects the fisheries sector, food systems, and regional economies. At the micro level, it affects household income, job security, consumption, and daily survival.

A good statistical system should be able to connect both levels. It should not only describe how much fish is produced or how many people work in fisheries, but also identify where risks are emerging and which groups may need support.

What the Indonesian case reveals

The Indonesian case shows that climate pressure on fisheries is not only an environmental problem. It is also a statistical visibility problem. Some risks are easy to observe locally, such as bad weather, failed fishing trips, or declining catches. But other risks are harder to detect nationally, including gradual changes in fishing activity, increasing operational costs, unstable income, or rising vulnerability among informal workers.

In this study, environmental variables were mapped across Indonesian waters for 2004, 2014, and 2024. Three indicators were used to represent changing marine conditions: sea surface temperature, salinity, and chlorophyll-a. Together, these variables capture ocean warming, changes in water characteristics, and marine productivity. Their spatial distribution shows that climate pressure is not uniform across Indonesian waters, making geospatial evidence essential for identifying where fisheries and coastal livelihoods may be most at risk.

The analysis shows a clear decline in fisheries potential between 2004 and 2024, indicating growing pressure on marine biodiversity. Looking ahead, sea temperatures are projected to increase by around 0.10–0.28°C within the next two to three decades. Under future climate scenarios, projected fisheries productivity in 2050 could decline substantially, with sharper losses under a high-emission pathway. These results on Figure 3 suggest that climate pressure is not only a future environmental concern, but already part of the changing conditions facing Indonesia's fisheries.

Spatial analysis further identifies two high-risk Fisheries Management Areas. The Arafura Sea faces high environmental vulnerability combined with intensive large-vessel activity, indicating a risk of ecological pressure and over-exploitation. The Java Sea, meanwhile, is highly vulnerable from a socio-economic perspective because

Data source	What it helps reveal
Labour statistics	Fisheries workers, informality, employment conditions
Household statistics	Welfare, poverty, consumption, food security
Vessel activity data	Fishing movement and operational intensity
Remote sensing/climate data	Environmental pressure and coastal exposure
Geospatial data	Location, accessibility, and spatial vulnerability

Table 1. Examples of data sources for assessing climate vulnerability in Indonesian fisheries.

Figure 2. An integrated data ecosystem for climate-sensitive fisheries statistics.

of the dominance of small-scale and artisanal fisheries. In this case, climate-induced fisheries depletion may directly affect many workers and households with limited adaptive capacity.

Integrated data can help reveal four important insights. First, climate pressure is spatially uneven. Not all coastal areas face the same level of risk. Some communities are more exposed because of geography, while others are more vulnerable because of limited infrastructure, weak market access, or dependence on small-scale fishing.

Second, fishing activity may respond to environmental change. Vessel movement, trip intensity, and fishing patterns can provide early signals that marine conditions are shifting. Third, livelihood vulnerability is socio-economic, not only ecological. The same environmental shock can have different consequences depending on household welfare, employment status, savings, alternative income sources, and access to services. Fourth, integrated evidence can support earlier policy response. When environmental signals, vessel activity, and household vulnerability are analyzed together, statistics can move beyond describing the past. They can help identify emerging risks before they become deeper crises.

Why this matters for official statistics

Climate change creates a new challenge for national statistical offices. Traditional statistical systems often measure population, employment, welfare, production, and prices as separate domains. But climate risk does not follow these boundaries.

A climate shock can begin as an environmental event, move through an economic sector, affect household welfare, and eventually become a social protection issue. If data systems remain fragmented, the full pathway may be missed. This does not mean replacing censuses and surveys. They remain the backbone of official evidence. But they can be strengthened by administrative data, remote sensing, vessel activity data, and geospatial layers.

In the climate era, the role of official statistics is not only to describe what has happened. It is also to help societies anticipate who is becoming vulnerable, where risks are emerging, and what kind of response is needed.

Conclusion

Climate change at sea is often invisible from land. But its consequences are not. They appear in uncertain fishing days, unstable household income, changing vessel movement, and the quiet anxiety of coastal communities facing a less predictable ocean. For Indonesia, understanding these risks requires a new kind of evidence system: one that connects the sea with the shore, the vessel with the household, and the climate signal with the human consequence.

By bringing together official statistics, geospatial evidence, vessel activity data, and environmental indicators, Indonesia can make hidden coastal risks visible. And once risks become visible, they can become actionable. In a warming world, fisheries statistics should not only count fish, vessels, or workers. They should help protect livelihoods.

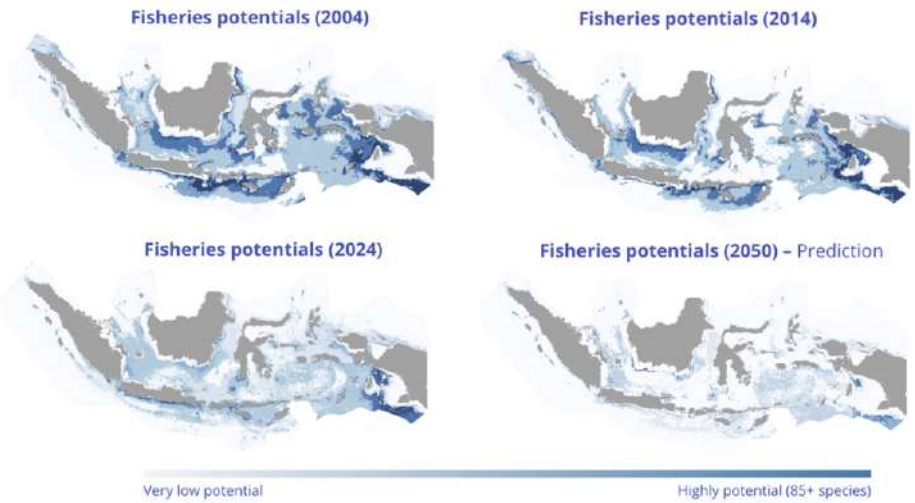

Figure 3. Spatial decline in fisheries potential across Indonesia from 2004 to 2024 and 2050, based on habitat suitability modeling.

COMMUNITY CORNER

ISI General Assembly 2026

The General Assembly is composed of all individual members of the ISI and holds the highest authority within the organisation. The **2026 ISI General Assembly** will take place **online** on **Wednesday, 1 July 2026, 14:00–15:30 CEST** (Zoom). All ISI members are cordially invited to attend. The Zoom link will be published on the ISI Communities platform and sent by email prior to the meeting. In accordance with item 6.3b of the Statutes, members are invited to propose agenda items by contacting ISI Board Secretary **Nurjahan Parent** (boardsecretary@isi-web.org) no later than **31 March 2026, 12:00 CEST**. The **Annual Report 2025** will be shared on the Communities platform on 1 June 2026. Members will vote online to approve the report in advance of the meeting; voting opens on **1 June** and closes on **24 June 2026**. Results will be announced during the General Assembly. More info and registration: communities.isi-web.org/isi-ga-2026.

ISI Service Awards 2025

The ISI Service Awards recognise members who have made outstanding contributions to the ISI and the international statistical community. Recipients are selected every two years by the ISI Executive Committee and honoured at the Awards Ceremony during the World Statistics Congress, where they receive a certificate.

2025 Recipients

Marianthi Markatou

Outstanding leadership as Co-Editor-in-Chief of the *International Statistical Review* (Mar 2022–Dec 2025).

Scott Holan

Outstanding leadership as Co-Editor-in-Chief of the *International Statistical Review* (Sep 2019–Dec 2023).

Ola Awad

Outstanding service as ISI Membership Elections Committee Chair.

Jogesh Babu

Outstanding service as ISI Astrostatistics Special Interest Group Chair.

Paulo Canas Rodrigues

Outstanding service as ISI Special Interest Group on Data Science Chair.

Jose Rosero Moncayo

Outstanding service as ISI Committee on Agricultural Statistics Chair.

Mara Sherlin Talento

Outstanding service as ISI Young Statisticians Committee Chair.

Adrian Röllin

Outstanding service as SPC Chair for the 65th ISI World Statistics Congress, The Hague 2025.

Oliver Chinganya

Outstanding service to ISI and the statistical community in Statistical Capacity Building developments and innovations.

Ksenija Dumičić

Leadership and commitment to ISI and her instrumental role in promoting the Committee on Women in Statistics.

Jūratė Petrauskienė

Outstanding service as SPC Co-Chair for the ISI–IAOS 2024 Conference in Mexico.

Manuel Mendoza

Outstanding service as SPC Co-Chair for the ISI–IAOS 2024 Conference in Mexico.

Other Awards & Recognitions

- **Bernoulli Society New Researcher Award (NRA) 2026** — Winners: **Morris Ang** (UC San Diego, random conformal geometry); **Serte Donderwinkel** (University of Groningen, random trees & graphs); **Mark Sellke** (Harvard University, probability, mathematical physics & theoretical CS). Honorable mentions: **Ahmed Bou-Rabee** (UPenn), **Sky Cao** (MIT), **Alexander J. Dunlap** (Duke).
- **David G. Kendall Award 2025 (BS/RSS)** — Winner: **Wei Qian** (University of Hong Kong), for exceptional contributions to random geometry, planar stochastic processes, and mathematical physics. Honorable mentions: **Samuel Johnston** (King's College London) and **Emma Horton** (University of Warwick).
- **JDSSV Best Paper Prize (inaugural)** — **Jakob Raymaekers** & **Peter Rousseeuw**, for *"Handling Cellwise Outliers by Sparse Regression and Robust Covariance"*.
- **William Sealy Gosset Award (ISBIS)** — **Katherine Bennett Ensor** (Rice University, ISI WSC 2025); **Nicholas Fisher** (University of Sydney, ISI WSC 2023).
- **TIES Abdel El-Shaarawi Early Investigator Award 2026** — **Paola Crippa** (University of Notre Dame), for outstanding contributions to environmental statistics through innovative integration of statistical methodology with atmospheric and environmental sciences.

Newly Elected ISI Members 2025-2026

Join us in welcoming the newest members of the ISI community.

- Karol P. Binkowski (*Australia*)
- Paul Pao-Yen Wu (*Australia*)
- Fabio Mariano Bayer (*Brazil*)
- Pedro Luiz Valls Pereira (*Brazil*)
- Rob Deardon (*Canada*)
- Shirin Golchi (*Canada*)
- N. K. Newlands (*Canada*)
- Tao Wang (*Canada*)
- Hao Mei (*China*)
- Yumou Qiu (*China*)
- Mengxin Yu (*China*)
- Andreas Artemiou (*Cyprus*)
- Peter J. Mannepalli (*India*)
- Praveen G. Sanka (*India*)
- U. D. Priyadarshi (*India*)
- Vaidyanathan VS (*India*)
- Sadia Khalil (*Pakistan*)
- M. R. M. V. Lourenço (*Portugal*)
- Pei-Fang Su (*Taiwan*)
- Sounak Chakroborty (*USA*)
- Paola Crippa (*USA*)
- Sujit Kumar Ghosh (*USA*)
- Jennifer Lynn Green (*USA*)
- Monnie McGee (*USA*)
- Wanli Qiao (*USA*)
- Michelle Shardell (*USA*)
- Amy Wagaman (*USA*)
- Lihu Xu (*USA*)
- Derek Scott Young (*USA*)
- Xiufan Yu (*USA*)
- Panpan Zhang (*USA*)
- Xin Zou (*USA*)

New ISI Special Interest Group: Indigenous Statistics

In February 2026, the ISI Executive approved the establishment of a new Special Interest Group (SIG) on Indigenous Statistics. The Management Committee will be jointly chaired by **Michele Connolly** *(USA)* and **Lisa Jackson-Pulver** *(Australia)*, both Indigenous women with distinguished records in the field. Membership will be open to all persons interested in the development of Indigenous statistics. Although several international agencies have expressed strong interest in Indigenous statistics, no single agency currently holds a leadership role in their development. The ISI and this SIG are well placed to fill that gap and take an influential role in shaping the field. Planned activities will be outlined at the upcoming ISI World Statistics Congress in Lusaka. For more information on ISI Special Interest Groups: https://isi-web.org/special-interest-groups

Upcoming Webinars

- **10 Jun 2026** — IASE Webinar: *Generation AI: A Finnish approach to K-9 AI Education.*
- **11 Jun 2026** — ISI Academy: *Measuring Wealth in Household Surveys in Low- And Middle-Income Countries.*
- **02 Jul 2026** — ISI Academy: *Recent Developments in Indigenous Statistics.*

66th ISI World Statistics Congress 11–15 July 2027 · Lusaka, Zambia

The **66th ISI World Statistics Congress** will take place from **11 to 15 July 2027** in *Lusaka, Zambia.* This leading global statistics event will bring together statisticians and data scientists from academia, official statistics, central banks, the private sector, and beyond to exchange knowledge, explore new developments, and build lasting professional connections.

Lusaka, a vibrant and fast-growing capital city at the heart of southern Africa, will provide an inspiring and memorable setting for learning, collaboration, and exchange. Working closely with our Zambian colleagues, we are preparing an inclusive and welcoming experience for all participants, held at the *Mulungushi International Conference Centre.*

The 2027 programme will feature insightful sessions on emerging trends and cutting-edge developments in statistics and data science, opportunities for hands-on learning, and dedicated spaces to connect and collaborate across sectors and regions. ISI and its Associations are made by, and not only for, their members — your participation and contribution help shape the future of our discipline.

Submissions are open. The system is accepting **Invited Paper Session (IPS) proposals** and **Contributed Paper Session (CPS) abstracts** for both talks and posters. Sessions are 100 minutes in length and may include 3–5 speakers, with up to 3 discussants or panellists. Mark your calendars and stay tuned for further details.

More info and submissions: https://www.isi-next.org/conferences/isi-wsc2027/

Calendar of Conferences & Workshops

- **14–20 Jun 2026** — 2026 Conference on Stochastic Processes and Their Applications (SPA 2026), *Ithaca, NY, USA.*
- **15–19 Jun 2026** — Conference on Small Area Estimation, Survey and Data Science (SAE 2026), *Bucharest, Romania.*
- **29 Jun–1 Jul 2026** — Data Science, Statistics & Visualization (DSSV 2026), *Trento, Italy.*
- **30 Jun–3 Jul 2026** — Computational and Applied Statistics for Data & AI (ICCSA 2026), *Braga, Portugal.*
- **8–10 Jul 2026** — 10th International Conference on Agricultural Statistics (ICASX 2026), *Krakow, Poland.*
- **12–17 Jul 2026** — 12th International Conference on Teaching Statistics (ICOTS 12), *Brisbane, Australia.*
- **12–16 Jul 2026** — 33rd International Biometric Conference, *Seoul, South Korea.*
- **20–21 Aug 2026** — 13th IFC Biennal Conference, *Basel, Switzerland.*
- **24–28 Aug 2026** — European Meeting of Statisticians (EMS), *Lugano, Switzerland.*
- **25–28 Aug 2026** — International Conference on Computational Statistics (COMPSTAT 2026), *Athens, Greece.*
- **07–10 Sep 2026** — RSS International Conference 2026, *Bournemouth, United Kingdom.*
- **18–20 Nov 2026** — 14th International Conference Use of R in Official Statistics (uRos2026), *Paris, France.*
- **07–11 Dec 2026** — LACSC-TIES-EnviBayes-EnvrASA 2026, *Mexico City, Mexico.*
- **09–11 Dec 2026** — IASC-ARS & SoAI 2026 Joint Conference, *Chiang Mai, Thailand.*

International Statistical Poster Competition 2026–2027

The Poster Competition has started in February 2026.

Participating countries can submit winning posters until 14.3.2027.

To sign up, please contact **islp.coordination@gmail.com**

jmp
STATISTICAL DISCOVERY

The ISI Mission

Established in 1885, the ISI is a non-profit, non-government organisation with members in over 150 countries. Our mission is to promote the understanding, development, and good practice of statistics worldwide by fostering knowledge, sharing best practices, and creating opportunities to network. Since 1947, the ISI has held Special Consultative Status with the UN Economic and Social Council, enabling participation in the UN Statistical Commission and related activities. The ISI network spans most national statistical offices, international organisations, and professional societies around the world.

Join the ISI

We welcome members who share our mission of *Statistical Science for a Better World*. The ISI unites statisticians and data scientists across sectors and regions, including students and early-career members. As a member, you gain recognition, community, up-to-date knowledge, discounts, and a growing international network. Membership runs from 1 January to 31 December.

Become a member

Existing or former members: log in to the Membership System with your email (reset password if needed) to renew, reinstate, or add memberships. For assistance: isimembership@isi-web.org.

ISI Journals

International Statistical Review

The flagship journal of the International Statistical Institute and its Associations, publishing broad-interest work in statistics and probability: reviews and expository pieces, statistical computing and graphics, statistics education, and applications across sectors.

Stat

An ISI–Wiley journal for rapid, high-quality dissemination of new theoretical, methodological, and applied results. Compact articles with fast decisions, supported by online materials (data, code, graphics, and media) across all areas of statistics and allied quantitative fields.

Executive Committee

President (2025–2027)
Fabrizio Ruggeri

President-Elect (2025–2027)
Nalini Ravishanker

Vice President (2023–2027)
Denise Britz do Nascimento Silva

Vice President (2023–2027)
Jessica Utts

Vice President (2025–2029)
Ayşe Bilgin

Vice President (2025–2029)
Matthew Shearing

Executive Director
Conchita Kleijweg

Council 2023–2027

Kaye E. Basford
Arup Bose
Miguel de Carvalho
Ying Chen
Daniel Frischemeier
Matthew Parry
E. Ashley Steel
Linda J. Young

Council 2025–2029

Olawale Ale
Cathy Chen
Rabaa Baniyas
G. Van Haldaren
Elisabetta Carfagna
V. Gonzalez-Lopez
M. Cetinkaya-Runderl
Arun Sinha

Association Presidents

Bernoulli Society
Nancy Reid

International Association for Official Statistics
Gary Dunnet

International Association for Statistical Computing
Paulo Canas Rodrigues

International Association for Statistical Education
Bruno de Sousa

International Association of Survey Statisticians
Partha Lahiri

International Society for Business and Industrial Statistics
Matilde Bini

The International Environmetrics Society
Nathaniel Newlands