

Methods note · Communities and governance

Publishing community maps without exposing lives

Ivalu Petersen · 5 October 2026

Fictional author and institute · Editorial demonstration, not peer reviewed

ABSTRACT

Public release changes what a community map can reveal, even when names are removed. This methods note compares NIST de-identification guidance with Indigenous data-governance and Canadian research ethics sources to examine identification, collective exposure, and inference from omissions. A wholly fictional seasonal-observation case contrasts exact points, displaced points, broad zones, and a non-spatial account. It shows why record counts, contributor counts, and complete release packages require separate scrutiny, and why suppressing one cell may not prevent reconstruction. The analysis favours a purpose-led, authorised derivative rather than automatic export of the source map. It provides neither an anonymisation guarantee nor a universal Indigenous checklist, and identifies changing linkage opportunities, review capacity, and irreversible distribution as important limits.

Keywords: community mapping · disclosure risk · data governance · spatial aggregation

A source-based editorial text, not a report of original Still North research. Examples and imagery are illustrative.

The publication question is narrower than the mapping question

A map made for a defined group can become a different object when it is searchable, downloadable, and combined with other records. Removing names does not undo that change. A route can suggest who travelled, a repeated observation can reveal a household's routine, and an apparently empty area can disclose where information was deliberately withheld. The question considered here is not whether community mapping is valuable. It is which public representation, if any, can serve an agreed purpose without exposing the relationships that made the source map meaningful.

This methods note compares disclosure controls described in NIST Special Publication 800-188 with the authority and collective-interest questions raised by Inuit Tapiriit Kanatami's research strategy and Canadian research ethics policy. It is a bounded synthesis, not a systematic review, an anonymisation certification, or a report of community consultation. The worked case is entirely fictional. No location, boundary, number, or publication preference represents a real community.

Separate three kinds of exposure

The first exposure is identification: connecting a record with a person or household. A coordinate, date, unusual activity, and public photograph can together identify someone even when none is a name. NIST distinguishes direct identifiers from quasi-identifiers and stresses that masking alone need not provide adequate de-identification. In a small population, a description that seems ordinary to an outside editor may be distinctive to a neighbour.

The second exposure concerns places and relationships rather than identifiable people. A layer may reveal a sensitive habitat, a culturally significant place, or a pattern of

land use. Aggregation that protects individual identities can still expose these things. Consequently, an individual privacy calculation cannot settle every publication decision. The CARE Principles explicitly situate Indigenous data within collective benefit, authority to control, responsibility, and ethics. They are not a licence for an outside institution to decide what a particular people should disclose.

The third exposure is inference about absence. Readers may interpret a blank cell as no activity, no concern, or no knowledge. In fact, it might indicate no participation, an agreement not to publish, or deliberate non-collection. A privacy-preserving map must therefore explain its coverage without disclosing why a particular household or place is missing. Sometimes this requires describing omissions for the product as a whole rather than attaching a revealing label to each gap.

Fictional case: a seasonal observations map

Imagine a project whose authorised internal archive contains 48 seasonal observations from 12 consenting contributors. The observations include precise coordinates and dates. The proposed public purpose is modest: explain that observation coverage differs between broad landscape zones. It is not to show travel routes or compare contributors. These invented counts illustrate design choices; they are not a minimum acceptable sample or a recommended disclosure rule.

The point map offers more detail than the stated purpose needs. Randomly shifting points is not automatically better. A shifted point may land in water, on another person's property, or close enough to its source to remain recognisable. Multiple independently shifted releases might also allow an observer to estimate the original location. Any displacement must be described as

Candidate product	What it preserves	What it can expose or distort
Unnamed points with exact dates	Fine spatial and temporal detail	Linkage to known journeys and rare activities
Points moved randomly within a fixed radius	Approximate spatial pattern	Residual linkage, implausible positions, and a false impression of measured location
Broad zones with seasonal count bands	Coarse coverage comparison	Small-group inference and loss of within-zone variation
A non-spatial coverage account	The explanation that coverage is uneven	Less geographic usefulness, but potentially fewer location disclosures

a transformation, not represented as measurement uncertainty in the underlying observation.

Suppose the broad-zone draft contains one zone with eight observations, all from a single contributor. An eight-record publication threshold would not protect that contributor merely because the count looks substantial. Repeated records and independent contributors are different denominators. Even contributor counts do not settle the issue if everyone in a zone shares a sensitive characteristic. The relevant question is what an informed reader could learn, not whether a familiar numeric threshold has been crossed.

Now suppose totals are released for the whole region and for all zones except one. Subtraction can reconstruct the omitted count. Suppressing one small cell while publishing every margin may therefore achieve little. The fictional editor would compare the complete release, including tables, captions, downloadable files, and earlier editions, rather than approve the image in isolation. The defensible result could be wider grouping, fewer totals, a different explanatory product, or no public map.

Define the threat model before choosing a technique

A threat model states what is being protected, from whom, with what plausible auxiliary information, and against which consequences. It need not assume a sophisticated attacker. Someone who knows a contributor's schedule may need only a public timestamp. A commercial user may combine downloadable layers at scale. A local reader may infer something an international researcher cannot. These are different disclosure pathways and call for different scrutiny.

NIST presents several sharing models, including public release, synthetic data, protected query interfaces, and non-public environments. That distinction is useful here because unrestricted download is not the only route to research usefulness. An authorised analyst might need precise locations under an agreement, while a public reader needs only the broad result. Access tiers should follow actual authority and purpose, not become a device for giving outside institutions preferential access to knowledge they cannot legitimately use.

Formal privacy methods can be appropriate for some statistical releases, but naming a technique does not establish its guarantee. Differential privacy, for example, requires a defined neighbouring-dataset relationship, a privacy budget, and control over composition across releases. Protection for one person's contribution does not by itself establish protection for a collectively significant

place. A small mapping project should not claim a formal guarantee merely because it adds random noise to coordinates or counts.

Publication is an authorised derivative, not a neutral export

The National Inuit Strategy on Research links Inuit governance with access, ownership, control, funding, and capacity. Its Inuit Nunangat scope matters. It must not be converted into a universal Indigenous publication checklist. In an actual project, the relevant authorities and knowledge holders would determine how their own protocols apply, including whether a public map should exist and who can review a derivative.

TCPS 2 Chapter 9 likewise distinguishes individual consent from community engagement and recognises that customs can govern what may be shared and with whom. Neither a participant signature nor institutional ethics approval should be portrayed as resolving every collective authority question. Conversely, community-level engagement does not erase an individual's choice not to participate. The publication record should preserve both distinctions without exposing private deliberations.

For the fictional case, a useful release record would identify the approved purpose, transformation, audience, review date, and responsible role. It would explain that public zones are not exact observation footprints and that count bands cannot support fine-scale comparisons. The sensitive source archive would retain the information needed for authorised work. The public package would not accidentally carry precise coordinates in hidden table columns, file properties, or a linked download.

Review also needs resources. Asking people to inspect a complex geospatial package without time, support, or usable tools does not create meaningful authority. The institution proposing publication should budget for the agreed review process and explain the difference between correcting the map, withholding material, and approving a scientific interpretation. Attendance at a workshop is not evidence that a whole community endorsed a release.

Limits and the smallest useful map

No technique makes a release permanently harmless. Auxiliary information changes, later datasets create new linkage opportunities, and copies can outlive the service that first distributed them. This note does not estimate re-identification probabilities, identify a legally sufficient anonymisation standard, or settle competing rights. Specialist disclosure assessment and locally applicable

governance remain necessary where consequences are material.

The useful implication is nevertheless concrete. Start with the claim the public needs to understand, then ask whether a map adds something essential. If it does, publish the least revealing authorised representation that still supports that claim, explain the lost detail, and assess the whole package across releases. A carefully limited map can remain informative. A decision not to map can be equally rigorous when the apparent precision would expose more than the research is entitled to share.

References

1. Garfinkel, S., Guttman, B., Near, J., Dajani, A., and Singer, P. (2023). De-Identifying Government Datasets: Techniques and Governance. NIST SP 800-188.
2. Inuit Tapiriit Kanatami (2018). National Inuit Strategy on Research.
3. Canadian research agencies (2022). TCPS 2, Chapter 9: Research Involving the First Nations, Inuit, and Métis Peoples of Canada.
4. Carroll, S. R., et al. (2020). The CARE Principles for Indigenous Data Governance. *Data Science Journal*, 19, 43.