

Methods note · Wildlife and ecology

What counts as an independent tundra plot?

Priyanka Das · 16 October 2026

Fictional author and institute · Editorial demonstration, not peer reviewed

ABSTRACT

Large numbers of tundra measurements need not provide large numbers of independent treatment tests. This methods note synthesises experimental-design principles, International Tundra Experiment guidance, and clustered spatial-sampling methods. It distinguishes measurement, quadrat, plot, block, and site, and separates spatial autocorrelation from the more fundamental problem of unreplicated treatment assignment. Fictional designs contrast intensive subsampling of one treated patch with replicated paired blocks, showing which level supports the treatment estimate. The discussion considers blocking, environmental side effects, hierarchical analysis, design-respecting resampling, regional generalisation, and the allocation of effort between breadth and detail. It offers no universal spacing or minimum replicate count. Its conclusion is a reporting practice that states sample sizes at each consequential level and matches every inferential claim to the independence the actual design supports.

Keywords: tundra plots · pseudoreplication · spatial autocorrelation · experimental units · blocked designs

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

Many measurements can belong to one experimental unit

A hundred vegetation measurements can provide detailed information about a small patch of tundra without providing a hundred independent tests of warming. Leaves share plants, quadrats share plots, plots share hillslopes, and nearby sites share weather and history. The analytical problem is not that related observations are useless. It is that their level of independence must match the question being tested. More subsamples cannot replicate a treatment applied only once.

This methods note examines spatial autocorrelation and pseudoreplication in tundra plot studies. It draws on Hurlbert's experimental-design analysis, the International Tundra Experiment's basic objectives and design guidance, and Olofsson and colleagues' treatment of clustered spatial samples. It offers hypothetical designs, not a reanalysis of a real experiment. Every plant response and numerical result below is invented to illustrate inference; none represents observations made by Still North.

Start with the unit that could receive a different treatment

An experimental unit is the unit to which a treatment is independently assigned. If separate warming chambers are assigned to plots, plots may be experimental units. Ten quadrats inside one chamber are usually subsamples of that unit. If an entire hillslope is warmed or otherwise treated as one intervention, plots within it do not independently replicate the hillslope treatment. The design, rather than the spreadsheet row, determines this hierarchy.

Hurlbert's definition of pseudoreplication covers tests of treatment effects where treatments are unreplicated or supposed replicates are not statistically independent. The

problem is a mismatch between the hypothesis and the error term used to test it. A dataset can therefore have careful measurement, correct arithmetic, and an invalid treatment test. The remedy begins with design and an explicit scope of inference, not with adding a more impressive statistical label afterwards.

Observational studies require similar care, although their causal claims differ. Ten plots in wet tundra and ten plots in dry tundra can describe those sampled locations. If all wet plots occupy one valley and all dry plots another, vegetation type is confounded with valley history. A model may describe the association, but cannot establish that moisture alone caused the difference. Measuring a covariate does not automatically remove every unmeasured site difference.

Spatial autocorrelation is a property to investigate

Nearby observations can resemble one another because they share soil, snow, drainage, grazing, disturbance, or clonal plants. Positive spatial autocorrelation reduces the amount of independent information in a cluster relative to equally many well-distributed observations. Its scale depends on the response and environment. There is no universal spacing distance that guarantees independence for every tundra variable.

The appropriate question is often whether unexplained residual variation remains spatially structured after relevant design terms and covariates are considered. A broad vegetation gradient in raw measurements is not identical to residual correlation. Mapping coordinates, examining distance relationships, and understanding the landscape can guide the analysis. Statistical diagnostics should complement ecological reasoning, not turn a weakly sampled

spatial pattern into a supposedly well-estimated correlation model.

Olofsson and colleagues discuss a related issue for clustered map-reference samples: clustering can reduce travel costs, but positive within-cluster correlation reduces precision and requires a design-appropriate uncertainty estimate. The same allocation trade-off appears in plot work. A compact cluster may be efficient for describing a local patch while being inefficient for estimating a landscape mean. The target determines whether the cost saving is worthwhile.

Worked example: replication at the wrong level

Fictional worked example. A study places one large warming structure over a tundra patch and selects one nearby untreated patch. Twenty quadrats are measured inside each. Suppose mean plant height differs by six centimetres. Treating the forty quadrats as forty independent treatment units would confuse subsampling with replication. There is one treated patch and one comparison patch, so treatment is inseparable from any pre-existing patch difference not otherwise resolved by the design.

Repeated visits provide a richer trajectory, but do not create new independently treated patches. Adding plant identities, thousands of leaves, or another year can improve description of within-patch change. It cannot by itself establish how the treatment would perform across other patches. A mixed model with a patch term cannot manufacture replicated patch-level treatment assignment where none exists. The study may still be a useful case study if its claim is narrowed honestly.

Consider a second, explicitly hypothetical design. Four spatial blocks each contain one randomly assigned warming plot and one control plot. Ten quadrats are sampled in each plot. The observed block-level warming-minus-control differences are +2, +4, -1, and +3 centimetres. Their mean is +2 centimetres. If the four block differences can reasonably be treated as independent, their sample standard deviation is about 2.16 centimetres and their mean's standard error about 1.08 centimetres.

There are eighty quadrat measurements, but only four paired block contrasts for that simple treatment comparison. The quadrats improve estimation of each plot mean; they do not turn the standard error into one based on forty independent warming effects. The calculation is intentionally simplified and does not establish statistical significance. Baseline information, temporal structure, small-sample uncertainty, and the actual randomisation scheme would determine the full analysis.

Blocking removes some noise without making sites interchangeable

Blocks can group plots with similar environmental conditions, then distribute treatments within each block. This helps separate treatment from an existing gradient when the arrangement and assignment are appropriate. If all warmed plots are uphill and all controls downhill, adding block labels afterwards does not remove that confounding. The physical arrangement and randomisation should be recorded before outcomes are examined.

The International Tundra Experiment guidance explicitly connects plant responses with measurements of the altered environment and encourages comparable siting conditions. That is important because a warming treatment can also affect moisture, nutrients, or snow conditions. Calling the intervention warming does not prove that temperature is the only changed pathway. Record the treatment's actual environmental effects before attributing all plant response to a single mechanism.

Blocking within one site does not create regional replication. A well-designed local experiment can estimate an effect for its sampled environment while remaining uncertain about other soils, climates, species assemblages, and grazing regimes. Multiple sites broaden the evidence only if their selection and analysis support that broader target. Several nearby sites sharing one disturbance history should not automatically count as independent regional climates.

Analysis should preserve the hierarchy

Begin with a diagram or table linking measurement, quadrat, plot, block, site, and visit. Identify treatment assignment at the correct level. For a simple balanced design, analysis of plot means or paired block contrasts may be transparent and adequate. More complicated nesting or repeated measurements can justify hierarchical or correlation-aware models. Choose the model because the design requires it, not because complexity itself guarantees validity.

Resampling must also respect the hierarchy. A bootstrap that independently resamples leaves from a small number of plants can produce a very narrow interval while ignoring plant-to-plant variation. Resampling whole experimental units, or using a justified hierarchical scheme, better matches the target. A permutation test must likewise follow the assignments that could actually have occurred. Arbitrarily shuffling treatment labels among subsamples can test an experiment that was never conducted.

Spatial models can represent dependence and improve uncertainty estimates when the data constrain their parameters. They cannot resolve perfect confounding between treatment and location. Nor should they absorb an ecological response into a spatial term without considering the causal question. Report sensitivity to reasonable correlation structures where it materially affects the inference, particularly with few blocks or sites.

Allocate effort between depth and breadth

Additional subsamples are valuable when plot means remain noisy or within-plot heterogeneity is itself the question. Once those means are adequately constrained, another independent experimental unit may improve the treatment estimate more than another quadrat inside an existing plot. Travel, permissions, disturbance, and maintenance limit what is possible, so the optimal allocation is not purely statistical. The constraints should shape the claimed target, not be hidden in an inflated sample count.

Location records must be sufficient for analysis without automatically disclosing sensitive sites. Appropriate governance can permit controlled access or derived distances where exact coordinates should not be public. Protecting

those data does not require abandoning the design record. Retain the information under the agreed authority so dependence, disturbance, and revisits can be evaluated responsibly.

Limits and a useful sample-size statement

This note does not offer one universal spatial model, plot spacing, or minimum number of replicates. Responses differ in correlation scale, and small samples can leave dependence poorly estimated. The worked contrasts assume comparable measurement and independent blocks for illustration, not as a claim that four blocks are sufficient for every ecological question.

A useful report states sample sizes at every consequential level and names the level supporting each inference. Distinguish an experiment on independently assigned plots from a detailed case study of one treated patch. Both can contribute evidence if described honestly. The practical advance is not making every observation independent. It is recognising what each observation adds, so the next measurement strengthens the question rather than merely enlarging the table.

References

1. Hurlbert, S. H. (1984). Pseudoreplication and the design of ecological field experiments. *Ecological Monographs*, 54, 187–211.
2. International Tundra Experiment. Basic hypotheses and objectives, including experimental design. ITEX manual, chapter 2; objectives approved in 1992.
3. Olofsson, P., Foody, G. M., Herold, M., Stehman, S. V., Woodcock, C. E., and Wulder, M. A. (2014). Good practices for estimating area and assessing accuracy of land change. *Remote Sensing of Environment*, 148, 42–57.