

Methods note · Land and freshwater

Designing an area-weighted tundra snow survey

Amina Okafor · 4 October 2026

Fictional author and institute · Editorial demonstration, not peer reviewed

ABSTRACT

Accurate snow cores do not guarantee an area-representative catchment estimate. This methods note synthesises snow-sampling evidence, Arctic snow reporting, and uncertainty principles to examine stratification and area weighting for tundra snow water equivalent. A fictional three-stratum catchment demonstrates how equal sample allocation can overrepresent deep drifts, while a small drift area can still contribute substantially to stored water. The note distinguishes weighting from representative selection and follows sampling, map-area, spatial-dependence, and survey-timing uncertainties. It considers how the next sample can improve the target estimate rather than merely enlarge a convenient route. No operational travel advice or real snow assessment is provided. The proposed reporting account includes terrain shares, class means, selection design, missing coverage, and the date represented, without treating storage as a runoff forecast.

Keywords: snow survey · area weighting · stratified sampling · snow water equivalent · sampling support

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

Which average does the route produce?

A snow survey can collect accurate measurements and still give the wrong answer to a catchment question. If most samples come from sheltered drifts because they are accessible, an unweighted average describes the survey route rather than the landscape. If difficult drifts are omitted, the opposite bias may occur. The number of readings is not a substitute for knowing how each reading relates to the area whose snow storage is being estimated.

This methods note examines stratification, area weighting, and sample allocation for a tundra snow-water-equivalent survey. It synthesises Fassnacht's field-sampling analysis, the NOAA terrestrial snow assessment, and standard uncertainty principles. It is not an operational travel or avalanche guide. All terrain areas, snow values, costs, and uncertainty calculations below are fictional illustrations. The target quantity is mean snow water equivalent over a defined catchment on a specified date, including its snow-free ground.

Fix the target before dividing the landscape

Snow water equivalent, or SWE, is a water-equivalent depth. Averaged over the catchment, it can be converted to stored water volume by multiplying by catchment area. The target is not necessarily the mean among snow-covered locations, the mean depth along a transect, or the largest winter storage. Those are different statistics. A survey made after melt begins may be useful, but should not be labelled peak storage without evidence about timing.

The NOAA snow assessment distinguishes extent, duration, and SWE and explains why sparse ground observations must be combined with broader products for Arctic-scale analysis. At catchment scale, the same support problem appears in miniature. A sensor footprint

is not an entire terrain class. A route crossing several classes is not automatically an area-representative sample of them.

Useful strata are mutually exclusive areas that together cover the target landscape and separate predictable snow differences. Exposed ridges, open lowlands, shrub terrain, and drift-prone hollows are possible candidates, not a universal legend. Boundaries should be determined from an explicit mapping rule rather than drawn around the observed snow values after sampling. A practical design balances meaningful snow contrasts against the difficulty of mapping and sampling too many tiny classes.

Area weighting corrects allocation, not selection bias

Within each stratum, estimate mean SWE from a defensible probability sample or an explicitly model-based design. Multiply each stratum mean by its fraction of catchment area, and sum. Samples can be allocated disproportionately to difficult or variable terrain without biasing this mean, provided the analysis uses the correct design weights and the within-stratum sample is representative. Giving every bottle or core equal weight is justified only by an appropriate selection design.

Weighting cannot repair arbitrary within-stratum selection. If a team samples only shallow edges of deep hollows, multiplying by the hollow area still gives a biased estimate. Nor can it create information for an entirely unsampled stratum. A model may predict missing terrain using other observations, but that contribution should be labelled modelled and carry the uncertainty of extrapolation. Inaccessible does not mean snow-free.

Stratum	Area share	Illustrative mean SWE	Contribution to catchment mean
Exposed terrain	60%	50 mm	30 mm
Shrub lowland	30%	150 mm	45 mm
Drift hollows	10%	350 mm	35 mm

Worked example: more drift samples, less catchment snow

Fictional worked example. A mapped catchment comprises 60% exposed terrain, 30% shrub lowland, and 10% drift hollows. Their estimated mean SWE values are 50, 150, and 350 millimetres. The area-weighted catchment mean is $0.60 \times 50 + 0.30 \times 150 + 0.10 \times 350 = 110$ millimetres. Equal numbers of samples in each class would instead produce an unweighted mean of about 183 millimetres.

The small drift area contributes almost one-third of the estimated water storage. It should not be discarded as an outlier. At the same time, collecting half the survey's cores there would not mean that drifts occupy half the catchment. Over a hypothetical catchment of 20 square kilometres, 110 millimetres corresponds to 2.2 million cubic metres of water. This is a stock estimate at the survey date, not an estimate of the volume that will necessarily pass the outlet.

Suppose the estimated standard errors of the three stratum means are 5, 10, and 30 millimetres. Ignoring area uncertainty and assuming independent stratum estimates, the catchment-mean standard error is the square root of $(0.60 \times 5)^2 + (0.30 \times 10)^2 + (0.10 \times 30)^2$, approximately 5.2 millimetres. The example makes a useful point: the largest within-class uncertainty does not necessarily dominate the final estimate, because area weights also matter.

Choose the next sample by information, not habit

Under a simple stratified random design with similar costs and negligible sampling fractions, a common variance-efficient allocation assigns more samples to strata with larger area multiplied by within-stratum standard deviation. This follows from the variance of the weighted mean. If travel or core extraction costs differ, cost must enter the allocation decision. A design intended to estimate each class precisely may reasonably differ from one optimised only for the catchment total.

In practice, initial estimates of variability may be weak. A small pilot can reveal whether an apparently uniform class contains contrasting drift and scour zones. The pilot should inform a documented design rather than become a reason to repeatedly sample whichever route is easiest. Minimum coverage of every relevant stratum may be more valuable than nominal mathematical optimality based on uncertain variance estimates.

Fassnacht's sampling analysis demonstrates differences across fine spatial scales and between ground measurements and remotely sensed estimates. It also shows why position accuracy and measurement support affect comparisons. A core and a mapped pixel can disagree because they represent different ground, not because one instrument failed. A valid comparison should align footprints or

aggregate multiple ground measurements to the mapped support.

Spatial dependence changes the value of repetition

Ten cores placed close together may describe one local patch very well while adding little information about the rest of a stratum. Their errors in estimating the stratum mean are not necessarily independent. A cluster design can reduce travel costs, but the uncertainty analysis must recognise clusters. Treating every core as an independent catchment-wide replicate can make the final interval far narrower than the design supports.

Repeated transects can be valuable for change detection, especially where location histories are maintained. However, fixed-route repeatability and area representation remain separate properties. A route that consistently misses late-lying snowbeds can document a stable local series while underestimating catchment storage. Adding a spatially distributed probability sample alongside established routes may reveal that difference without abandoning the historical observations.

Map uncertainty and seasonality remain in the account

Area shares are estimated quantities. If shrub lowland is confused with exposed terrain, the weights themselves become uncertain. In the worked example, transferring five percentage points of area from exposed terrain to drift hollows raises the catchment mean by 15 millimetres. This sensitivity is larger than the illustrative sampling standard error. A more precise core balance would not address that source of uncertainty.

NIST's guidance supports combining uncertainty through the sensitivities of a calculation and including covariance. Area shares sum to one, so their uncertainties are constrained and cannot all vary independently. Measurement bias can also be shared across strata, for example through a common weighing error. Separate the uncertainty of individual cores, the sampling uncertainty of class means, and the uncertainty of mapped weights before combining them.

A multi-day survey has a temporal footprint as well as a spatial one. Melt, snowfall, and wind redistribution can alter storage while the team crosses the catchment. If exposed terrain is sampled first and drifts last, a weather change may be confounded with terrain class. Record dates and conditions, use repeat points where appropriate, and avoid calling the result a single-date map unless a defensible temporal adjustment supports that interpretation.

Limits and a more useful survey summary

This note assumes that field methods can measure SWE adequately at the selected sites. Core loss, vegetation, ice layers, and inaccessible depths can violate that assumption.

tion and need method-specific evaluation. It also does not guarantee that a chosen stratification captures the main sources of variability. Class boundaries and the sampling design should be revised when evidence shows that their original rationale no longer holds.

The useful survey summary is therefore not only an average and a sample count. It includes the target boundary and date, area shares, class means, selection design, missing terrain, and uncertainty contributions. Such an account can distinguish a genuinely snow-rich catchment from a route rich in drifts. That distinction gives the next survey a practical purpose: measure the places that most improve the answer, rather than merely repeat the largest number of readings.

References

1. Fassnacht, S. R. (2021). A call for more snow sampling. *Geosciences*, 11, 435. DOI: 10.3390/geosciences11110435.
2. Mudryk, L. R., and colleagues (2024). Terrestrial snow cover. NOAA *Arctic Report Card 2024*.
3. National Institute of Standards and Technology. Technical Note 1297, section 5: Combined standard uncertainty.