Catalogue no. 12-001-X
ISSN 1492-0921

Survey Methodology

Master samples with optimized panels

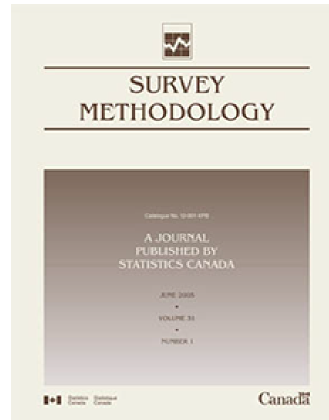

by Anton Grafström

Release date: June 29, 2026

Statistics Canada   Statistique Canada

Canada

How to obtain more information

For information about this product or the wide range of services and data available from Statistics Canada, visit our website, www.statcan.gc.ca.

You can also contact us by

Email at infostats@statcan.gc.ca

Telephone, from Monday to Friday, 8:30 a.m. to 4:30 p.m., at the following numbers:

• Statistical Information Service	1-800-263-1136
• National telecommunications device for the hearing impaired	1-800-363-7629
• Fax line	1-514-283-9350

Standards of service to the public

Statistics Canada is committed to serving its clients in a prompt, reliable and courteous manner. To this end, the Agency has developed standards of service which its employees observe in serving its clients. To obtain a copy of these service standards, please contact Statistics Canada toll-free at 1-800-263-1136. The service standards are also published on www.statcan.gc.ca under "Contact us" > "Standards of service to the public."

Note of appreciation

Canada owes the success of its statistical system to a long-standing partnership between Statistics Canada, the citizens of Canada, its businesses, governments and other institutions. Accurate and timely statistical information could not be produced without their continued co-operation and goodwill.

Published by authority of the Minister responsible for Statistics Canada

An HTML version is also available.

Cette publication est aussi disponible en français.

Survey Methodology, June 2026
Vol. 52, No. 1, pp. 173-185
Statistics Canada, Catalogue No. 12-001-X

Master samples with optimized panels

Anton Grafström[1]

Abstract

We introduce a general framework for constructing master samples that preserve desirable design properties across panels. The core procedure is to order an initial probability sample. Since the final sequence must be robust to a uniform random rotation, we define and minimize an objective that aggregates panel-level performance across all possible circular panels. A final random rotation is applied to ensure design validity. The framework is flexible with respect to the choice of design criteria, such as spatial balance or marginal balance, and can be implemented efficiently using simulated annealing to obtain high-quality approximate solutions. By construction, the approach supports both positive and negative sample coordination for spatially balanced, marginally balanced, and doubly balanced samples. The method's versatility is demonstrated through three applications: constructing a master sample with spatially balanced panels, marginally balanced panels, and doubly balanced panels.

Key Words: Balanced sampling; Sample coordination; Sampling; Simulated annealing; Spatial balance.

1. Introduction

A master sample is a single carefully selected probability sample of the population that can be subdivided, reused, or coordinated to meet a variety of survey objectives over time (Van Dam-Bates, Gansell and Robertson, 2018; Larsen, Olsen and Stevens, 2008). By drawing subsamples or panels from a master sample, survey practitioners can achieve efficiencies that would be difficult or impossible to obtain if each survey were designed independently.

There are many practical motivations for using master samples. One common application is in the construction of rotating panel designs, where successive survey waves are based on subsets of the master sample. This makes it possible to balance the competing goals of estimating cross-sectional quantities and measuring change over time. Another use case is general sample coordination, in which panels are chosen to be either non-overlapping (negative coordination) or overlapping (positive coordination) in order to manage respondent burden, cost efficiency, or the precision of estimates in repeated surveys. Master samples also allow for flexible sample sizes: since the underlying sample is fixed, subsamples of varying sizes can be drawn to accommodate changing resource constraints or shifting survey objectives.

Together, these advantages make master samples a highly versatile component of modern survey methodology. At the same time, the statistical and computational tools for constructing master samples remain rather limited. Various methods are available to select a single high-quality initial sample, such as the local pivotal method (Grafström, Lundström and Schelin, 2012), spatially correlated Poisson sampling (Grafström, 2012), locally correlated Poisson sampling (Prentius, 2024), the cube method (Deville and Tillé, 2004) and doubly balanced sampling (Grafström and Tillé, 2013). Implementations of many of these foundational methods are available in the `BalancedSampling` R package (Grafström and Prentius, 2024). Recent

1. Anton Grafström, Department of Forest Resource Management, Swedish University of Agricultural Sciences, Umeå, Sweden. E-mail: anton.grafstrom@slu.se.

advances have been made in equal-probability master samples (Robertson, Price and Reale, 2024), with extensions to balanced and doubly balanced samples (Robertson, Price, Reale and Davies, 2026). Moreover, highly effective methods have been developed to directly construct spatio-temporal samples with spatial and temporal spreading properties (Eustache, Jauslin and Tillé, 2022). However, a general framework for transforming a single high-quality probability sample into a master sample whose ordering is optimized for arbitrary design criteria has been lacking. The framework introduced in this paper addresses this gap.

We apply Simulated Annealing (SA) to optimize the ordering of the master sample, thereby improving the quality of all induced panels. Importantly, our framework performs this optimization while strictly maintaining the design-based inclusion probabilities of the subsamples. This contrasts with other applications of SA in sampling, such as optimizing allocation and stratification (O'Luing, Prestwich and Tarim, 2022), or optimizing a single sample configuration (see e.g. Lark, 2016) in a way that does not constitute probability sampling.

The remainder of this paper is structured as follows. In Section 2, we formally introduce our general framework. In Section 3, we provide a detailed description of the optimization algorithm and the modular objective functions used to achieve different design properties. Section 4 presents three examples that demonstrate the possibilities of the framework. We conclude with a discussion in Section 5.

2. Master sample creation

Our proposed procedure for creating a high-performance master sample is a multi-stage procedure. It begins with a high-quality initial sample, finds a good ordering of that sample, and concludes with a randomization step that guarantees valid probabilities for subsamples.

2.1 Population and initial sample selection

Let $U = \{1, \ldots, N\}$ be a finite population of size N. The goal is to create a master sample of size $n < N$.

The first step is to select an initial fixed-size sample $S_n \subset U$ of size n. This sample must be drawn using a method that ensures a desired design property. For example, to create a master sample with spatially balanced panels, S_n must be selected using a spatially balanced sampling design.

Let the first-order inclusion probabilities for this initial sample be denoted by $\pi_i = \Pr(i \in S_n)$, for all $i \in U$. These probabilities are chosen by the survey designer, satisfy $\sum_{i \in U} \pi_i = n$, and may be proportional to an auxiliary variable. Once this initial sample S_n is selected, it becomes the input for the ordering stage.

2.2 Ordering of the initial sample

Although the set S_n may have the desired property as a whole, an arbitrary or random ordering does not guarantee that its subsequences will also share this property. The second stage of our framework is therefore to find a high-quality ordering of the units in S_n.

Let $S^{\circ} = (s_1^{\circ}, s_2^{\circ}, \ldots, s_n^{\circ})$ be a permutation of the units in S_n. The goal is to find an ordering S° that is close to optimal with respect to some quantifiable criteria, typically a form of robustness of its subsamples. This can be achieved using any suitable optimization algorithm. In Section 3, we describe in detail a method for finding such an ordering using Simulated Annealing. For now, we consider S° to be the final ordered sequence that is the result of this search.

2.3 Random rotation and final inclusion probabilities

The third and final stage is a randomization step that provides the necessary probability structure for the master sample. We apply a uniform circular random shift to the optimized sequence S°. Let J be a random starting position drawn uniformly from $\{1, \ldots, n\}$. Define the circular index function

$$\iota(j) = ((j-1) \bmod n) + 1, \qquad j \in \mathbb{Z}.$$

The final master sample, denoted $S^{\star}$, is then defined as the circular rotation

$$S^{\star} = (s_{\iota(J)}^{\circ}, s_{\iota(J+1)}^{\circ}, \ldots, s_{\iota(J+n-1)}^{\circ}).$$

Any subsample, such as a panel, is then taken as a contiguous block of a desired length k from this final sequence $S^{\star}$. This procedure leads to the following general theorem for the inclusion probabilities of any such subsample.

Theorem 1. (Subsample inclusion probabilities)**.** *Let S_k' be a subsample taken as a contiguous block of length $k \le n$ from the randomly-rotated sequence $S^{\star}$. For any unit $i \in U$, the first-order inclusion probability is given by:*

$$\pi_{i,k} = \Pr(i \in S_k') = \frac{k}{n} \pi_i.$$

Proof. The proof relies on the exchangeability induced by the random rotation. Let $I_i = \mathbf{1}(i \in S_n)$ be the indicator that the unit i is in the initial sample. By the design of the initial sample, $\Pr(I_i = 1) = \pi_i$. Uniform random rotation ensures that, conditional on the event $i \in S_n$, unit i has a $1/n$ chance of landing in any of the n positions in the final sequence $S^{\star}$. The probability of landing within a specific contiguous "window" of size k is therefore exactly k/n. By the law of total probability, the unconditional inclusion probability is the product of these two probabilities: $\pi_{i,k} = \Pr(i \in S_k' \mid i \in S_n) \times \Pr(i \in S_n) = (k/n)\,\pi_i$.

This theorem is powerful because it holds for *any* ordering S°. It successfully decouples the problem of satisfying the probability constraints from the problem of achieving, for example, high spatial quality. This allows us to use a global optimization search for the ordering stage, confident that the final random rotation will make the resulting subsamples probabilistically valid. The next section describes our method for finding a close to optimal ordering.

3. Optimization of the master sample ordering

Once the initial probability sample S_n has been selected, the only remaining degree of freedom in the construction of the master sample is the *order* of its n units. This ordering fully determines all possible subsamples that may be obtained after the random circular rotation described in Section 2. Consequently, the problem of constructing a high-quality master sample reduces to choosing an ordering of S_n such that all induced panels perform well with respect to a chosen design criterion.

3.1 The ordering problem

For a given panel size $k \leq n,$ the ordering S° induces n circular panels. Using the circular index function $\iota(\cdot)$ defined in Section 2, the panel of size k starting at position j is

$$S'_{k,j} = (s^{\circ}_{\iota(j)}, s^{\circ}_{\iota(j+1)}, \ldots, s^{\circ}_{\iota(j+k-1)}), \qquad j = 1, \ldots, n.$$

Because the final master sample is obtained by applying a uniform random circular rotation, every one of these n panels is equally likely to be selected. A desirable ordering is therefore one for which *all* circular panels exhibit good design properties. This motivates an objective function that aggregates panel-level performance across all rotations and across all panel sizes of interest.

3.2 Objective function

Let $Q(S'_{k,j})$ denote a nonnegative quality function that measures how well a single panel $S'_{k,j}$ satisfies a desired design criterion. Lower values of Q correspond to better performance. The choice of Q depends on the type of quality that we want the resulting panels to have. For example, we can target spatial balance, marginal balance, or a combination of both.

- *For spatial balance:* A sample is considered spatially balanced if its units are well-spread across the space of auxiliary variables. There exist different ways to measure spatial balance, such as using Voronoi polytopes (Stevens and Olsen, 2004), a measure based on Moran's I index (Tillé, Dickson, Espa and Giuliani, 2018) and a local balance measure (Prentius and Grafström, 2024). Here we use the measure based on Voronoi polytopes, which quantifies the variation in inclusion probability mass within the polytopes generated by the sample units. There are two possibilities for optimization. We can either aim to make each subsample as well-spread as possible within the full population (the unconditional measure), or treat this as a problem conditioned on the realized master sample, aiming for spread within the master sample (the conditional measure). In both cases, the general form of the measure is

$$Q_{\text{spread}}(S'_{k,j}) = \frac{1}{k} \sum_{i \in S'_{k,j}} (v_i - 1)^2,$$

where v_i is the total probability mass in the Voronoi polytope V_i that unit i creates. For the unconditional measure, V_i contains all population units closer to unit i (in the auxiliary space) than to any other unit in the panel $S'_{k,j}$, and v_i is computed with unconditional panel inclusion probabilities as $v_i = \sum_{u \in V_i} \pi_{u,k}$. For the computationally efficient conditional measure used in our optimization, V_i contains all units in the master sample S_n closer to i than to any other unit in the panel $S'_{k,j}$, and v_i is computed with the equal conditional probabilities as $v_i = \sum_{u \in V_i} (k / n) = (k / n) |V_i|$. A lower value of Q_{spread} indicates a higher degree of spatial balance.

- *For marginal balance:* A sample is (marginally) balanced on a set of auxiliary variables if the Horvitz-Thompson estimates of the population totals agree with the known population totals for these variables. Often, it is only possible to achieve an approximate balance for which the cube method (Deville and Tillé, 2004) provides a solution. The quality function $Q_{\text{balance}}(S'_{k,j})$ measures the deviation from perfect balance on a set of P auxiliary variables with known and nonzero population totals. We use the mean squared relative error:

$$Q_{\text{balance}}(S'_{k,j}) = \frac{1}{P} \sum_{p=1}^{P} \left(\frac{\sum_{i \in S'_{k,j}} \frac{z_{ip}}{\pi_{i,k}} - Z_p}{Z_p} \right)^2,$$

where $Z_p = \sum_{i \in U} z_{ip}$ is the known population total for variable p, and $\pi_{i,k}$ is the panel inclusion probability from Theorem 1. A lower value indicates better balance.

- *For double balance:* It is also possible to select doubly balanced samples which are spread (spatially balanced) on a set of auxiliary variables and approximately marginally balanced on a potentially different set of auxiliary variables. Grafström and Tillé (2013) introduced the local cube method which combines the local pivotal method with the cube method and achieves double balance. For doubly balanced problems, the quality is a composite of the two individual quality functions, which requires addressing their different numerical scales. The general form of the quality is a weighted sum:

$$Q_{\text{double}}(S'_{k,j}) = w Q_{\text{balance}}(S'_{k,j}) + (1 - w)\, Q_{\text{spread}}(S'_{k,j}) \tag{3.1}$$

where the preference weight $w \in [0, 1]$ allows a user to prioritize one objective over the other.

Let K denote the set of panel sizes that are relevant for the intended application. For a given ordering S°, we define the overall objective function as

$$O_d(S^\circ) = \frac{1}{|K|} \sum_{k \in K} \frac{1}{n} \sum_{j=1}^{n} Q_d(S'_{k,j}), \tag{3.2}$$

where $d \in$ {spread, balance, double} represents the quality to optimize. For a fixed panel size k, the inner average represents the expected quality of a panel drawn from the ordering S° under a uniform random circular rotation. The outer average assigns an equal weight to each target panel size. Thus, $O_d(S^\circ)$ can be interpreted as the expected panel-level quality induced by the ordering S°, averaged over all planned panel sizes and all possible rotations.

Let Ω denote the set of all $n!$ permutations of the initial sample S_n. We seek to find the optimal sequence that minimizes the objective function (3.2), i.e.

$$S_{\text{opt}}^{\circ} \in \underset{S^{\circ} \in \Omega}{\arg\min} \quad O_d(S^{\circ}).$$

Because the size of Ω is $n!$, exact optimization is computationally infeasible for realistic master sample sizes. Thus, the term “optimized” refers to approximate minimization of the objective function (3.2), rather than to the attainment of a guaranteed global optimum.

Remark 1. (Choosing the weight for double balance)**.** *Choosing an appropriate w can be challenging. For the demonstrations in this paper, we aim for a compromise where both contribute equally to the total cost in the initial state. We achieve this by calculating a data-driven weight w_{auto}, based on the scores of the initial random sequence. Let S° be the initial sequence, then*

$$w_{\text{auto}} = \frac{O_{\text{spread}}(S^{\circ})}{O_{\text{balance}}(S^{\circ}) + O_{\text{spread}}(S^{\circ})}.$$

This parameter-free approach provides a default for doubly balanced optimization. However, we emphasize that practitioners can manually set w in (3.1) to fine-tune the result or deliberately target solutions that favor balance or spread.

3.3 Simulated annealing

To find a master sample sequence that is close to optimal for all planned subsamples, we employ Simulated Annealing (SA). SA is a powerful metaheuristic search algorithm designed to find a near-optimal solution in a large and complex combinatorial solution space (Kirkpatrick, Gelatt and Vecchi, 1983). It is particularly well-suited for ordering problems and can effectively escape poor local optima to find a high-quality final solution.

The algorithm works by starting with an initial candidate sequence and iteratively proposing structured random modifications. Moves that decrease the sequence’s overall cost are always accepted, while moves that increase the cost are accepted with a probability that decreases over time, mimicking the metallurgical process of annealing. Our implementation consists of three core components: the state space and neighbor generation mechanism, the objective function that defines the cost of a sequence, and the annealing process itself.

A “state” in our system is a complete single ordering of the n units in the initial sample S_n. This is represented as a permutation vector $S^{\circ} = (s_1^{\circ}, s_2^{\circ}, \ldots, s_n^{\circ})$. The state space Ω is the set of all $n!$ possible permutations.

To move from a current ordering $S_{\text{current}}^{\circ}$ to a neighboring ordering, we use a stochastic swap mechanism that targets poorly performing panels while retaining full connectivity of the permutation space.

At each iteration, a panel size k is selected uniformly at random from the set K. With probability $q \in (0,1)$, the worst performing circular panel of size k is identified, and a position j_1 is selected uniformly at random from the indices belonging to this panel. With the remaining probability $1-q$, the position j_1 is selected uniformly at random from the set $\{1, \ldots, n\}$. Then, a second position j_2 is selected, uniformly at random, from $\{1, \ldots, n\} \setminus \{j_1\}$. The two units at positions j_1 and j_2 are then swapped to produce the ordering $S^{\circ}_{\text{candidate}}$.

With this targeted move (we use $q = 0.9$), most of the proposed swaps are a direct attempt to fix a known weakness of the current solution. By assigning a strictly positive probability of selecting both swap positions uniformly at random, every transposition of two positions can occur with positive probability, ensuring that the full permutation space can be explored.

The annealing algorithm proceeds as follows. Starting with a random initial sequence, it iterates for a fixed number of steps. In each iteration t:

1. A new candidate sequence $S^{\circ}_{\text{candidate}}$ is generated from the current sequence $S^{\circ}_{\text{current}}$ using the targeted swap described previously.
2. The cost change, $\Delta O_d = O_d(S^{\circ}_{\text{candidate}}) - O_d(S^{\circ}_{\text{current}})$, is calculated.
3. The new sequence is accepted as the next current sequence based on the Metropolis criterion:

$$\text{Accept } S^{\circ}_{\text{candidate}} \text{ if } \Delta O_d < 0 \text{ or if Uniform}(0,1) < e^{-\Delta O_d / T_t},$$

where T_t is the temperature at iteration t.

The temperature T_t is gradually reduced according to a geometric cooling schedule $T_{t+1} = \alpha T_t$, where the cooling rate α is a constant very close to 1 (e.g., 0.999). When the temperature is high, the algorithm readily accepts worse solutions, allowing it to explore the state space broadly. As the temperature cools, the probability of accepting worse solutions decreases, forcing the algorithm to converge on a high-quality local optimum. To further enhance the search, we employ a reheating mechanism: if no improvement to the best-found solution is made for a large number of iterations, the temperature is reset to a higher value to re-energize the search from its current position. The algorithm terminates after a fixed number of iterations, and the best sequence ever encountered is returned as the final optimized ordering.

4. A demonstration: Optimizing for different objectives

To demonstrate the capabilities of our framework, we apply it to three design goals: creating a master sample that is optimized for spatial balance, one that is optimized for marginal balance of auxiliary variables, and one that is optimized for double balance.

4.1 Example 1: Optimizing for spatial balance

This example mimics the selection of a temporary 5-year sample in a stratum (region 3) of the Swedish National Forest Inventory (NFI), see Grafström, Zhao, Nylander and Petersson (2017). To enable the use of auxiliary information in the continuous population setting of the NFI, sampling is conducted in two steps. First, a large sample of size $N = 100{,}000$ is selected using a uniform sampling intensity within the stratum. For the selected sampling units (each unit is a set of circular plots), auxiliary variables are extracted. This large sample then serves as the sampling frame for the actual 5-year sample of size $n = 360$, which is selected using the local pivotal method to ensure spatial balance.

The 5-year sample must then be split into five non-overlapping yearly samples, each of size $k = 72$, where one sample is inventoried each year of the 5-year cycle. To improve yearly estimates, it is important that each of the five yearly samples is spatially balanced. Although the NFI routinely uses several auxiliary variables, here we only simulate spatial coordinates.

Thus, we generated a population of size $N = 100{,}000$ with uniformly distributed spatial coordinates as auxiliary variables. The master sample of size $n = 360$ was then selected using the local pivotal method. The simulated annealing (SA) algorithm was run for 100,000 iterations, with an initial temperature of $T_0 = 0.03$, corresponding to approximately 10% of the initial cost. This results in an initial probability of about 0.37 for accepting a move that increases the objective function by 10%. The cooling rate was set to $\alpha = 0.999$, and reheating reset the temperature to half the initial value after 10,000 consecutive non-improving iterations.

The initial ordering was a simple random permutation of the master sample. After running SA, the optimized sequence achieved a worst-case spatial balance measure of 0.082 for panels of size $k = 72$, evaluated with respect to the full population auxiliary information. Since the master sample contains 360 possible consecutive panels of size 72, optimization ensures that every such panel is highly spatially balanced. This result also validates the use of the conditional measure of spatial balance for optimization in large populations and large master samples. The spatial balance performance of the panels under optimized and random orderings is summarized in Table 4.1.

Note that the best-performing panel among all panels in the random ordering has a spatial balance measure of 0.137, which is worse than the worst-performing panel in the optimized ordering. Thus, a uniform random starting position can safely be selected in the optimized sequence, and the five consecutive non-overlapping yearly panels of size $k = 72$ can be extracted. Based on Table 4.1, each of these panels is guaranteed to exhibit a high degree of spatial balance. The total run-time for this example, including optimization and post-evaluation of spatial balance against the population, was less than seven minutes.

Table 4.1
Spatial balance performance across all 360 consecutive panels of size $k = 72$ from a master sample of size $n = 360$ (lower values indicate better spatial balance)

Method	min(Q_{spread})	mean(Q_{spread})	max(Q_{spread})
Original (Random)	0.137	0.271	0.519
Optimized (Simulated Annealing (SA))	0.036	0.055	0.082

4.2 Example 2: Optimizing for marginal balance

In some surveys, the primary goal is not spatial spread, but ensuring that subsamples are well balanced on a set of known auxiliary variables. Our framework can be directly applied to optimize a master sample for this objective.

We created a population of size N = 100,000 with two auxiliary variables $z_1 \sim N(10, 2^2)$ and $z_2 \sim N(100, 25^2)$. The goal was to create a master sample of size n = 1,000 with equal inclusion probabilities. The inclusion probabilities themselves were also included as an auxiliary variable to ensure that the fixed sample size property was respected.

First, a high-quality initial sample was selected using the cube method to ensure that it was well balanced on the population totals of the auxiliary variables. The quality of this initial sample was excellent, with a mean squared relative error of only 4.45×10^{-8}.

The critical step is the ordering of these n = 1,000 selected units. To provide a baseline, we considered a simple random permutation of the units, which ensures valid probabilities but does not optimize for balance in its subsamples. To find a robustly balanced sequence among the 1,000! possible orderings, we used the SA algorithm. The objective function was (3.2), which seeks to minimize the mean of the imbalance for a set of critical sizes $K = \{100, 200\}$. We ran the SA algorithm for 100,000 iterations. The initial temperature was T_0 = 0.00003 (roughly 10% of the initial cost). The cooling rate was α = 0.999 and reheating was set to be done after 10,000 non-improving iterations.

The results, shown in Table 4.2, demonstrate a substantial improvement due to optimization of the ordering. For panels of size k = 200, the worst-case mean squared relative error, measured by $\max(Q_{\text{balance}})$, is reduced by approximately 99.2% compared to a random ordering. For panels of size k = 100, the corresponding reduction is approximately 98.7%. These results illustrate the ability of the proposed optimization framework to identify an ordering that robustly preserves the balance property across all panels and for multiple panel sizes. The total run-time for this example, including both optimization and post-evaluation of balance, was less than two minutes.

Table 4.2
Marginal balance performance across all panels from a master sample of size n = 1,000, measured by the balance index Q_{balance} (lower values indicate better balance)

Panel Size (k)	Method	mean(Q_{balance})	max(Q_{balance})
100	Original (Random)	3.73×10^{-4}	2.39×10^{-3}
100	Optimized (Simulated Annealing (SA))	6.63×10^{-6}	3.07×10^{-5}
200	Original (Random)	1.90×10^{-4}	1.10×10^{-3}
200	Optimized (Simulated Annealing (SA))	2.03×10^{-6}	9.21×10^{-6}

4.3 Example 3: Optimizing for double balance

Here, we generate a spatial population of N = 100,000 units by placing points independently and uniformly on the unit square. Two auxiliary variables were generated as $z_1 \sim N(10, 2^2)$ and $z_2 \sim N(100, 25^2)$.

We then selected a doubly balanced master sample of size $n = 500$, using the local cube method with equal inclusion probabilities. This sample is well spread in the coordinates in the unit square, balanced on the inclusion probabilities, and approximately balanced on the two auxiliary variables. Balancing on the inclusion probabilities ensures a fixed sample size. We then used the SA-algorithm to optimize this master sample for panels of sizes $K = \{50,100\}$, and as a baseline we compare against the initial simple random permutation. We used 100,000 iterations, an initial temperature of $T_0 = 0.0001$ (10% of initial cost), a cooling rate of $\alpha = 0.999$ and reheating was set to be performed after 10,000 non-improving iterations.

The results in Table 4.3 show a substantial improvement for both criteria. For panels of size 100, the worst-case mean squared relative error is reduced by approximately 96%, while the worst-case spread is reduced by about 69%. For panels of size 50, the worst-case mean squared relative error is reduced by approximately 96%, and the worst-case spread is reduced by about 77%. The total run-time for this example, including optimization and post evaluation of balance and spatial balance, was 35 minutes.

Table 4.3
Worst-case performance of balance and spread across all panels (lower is better)

Panel Size (k)	Method	$\max(Q_{\text{balance}})$	$\max(Q_{\text{spread}})$
50	Original (Random)	2.97×10^{-3}	0.592
50	Optimized (Simulated Annealing (SA))	1.15×10^{-4}	0.139
100	Original (Random)	1.02×10^{-3}	0.392
100	Optimized (Simulated Annealing (SA))	4.11×10^{-5}	0.122

5. Discussion

We have introduced a general framework that can be used to produce master samples. This framework allows for the creation of master samples selected without replacement with probability proportional to size, where subsamples of any planned size have both valid probabilities and a high degree of a desired design property, such as spatial balance, marginal balance, or both. The flexibility to tailor the objective function makes the method very general. For example, to create a highly efficient system of overlapping panels, we include both the panel size and the overlap size in the objective function. This guarantees, for example, that both the panels and their overlap are well balanced. Thus, it supports both positive and negative sample coordination.

The primary limitation is the computational cost, which depends on the size of the master sample, the complexity of the objective function, and the efficiency of the search algorithm for the optimal order. However, the construction of a master sample is generally a one-time operation, where the computational cost of optimization is likely to be a small fraction of the total survey cost. The examples in this paper took between 1 and 35 minutes to run. For applications requiring the highest degree of robustness and quality, this small upfront investment yields a superior and more reliable result.

The proposed framework is particularly well-suited for stratified environmental monitoring programs, such as the Swedish NFI or the European Union’s Land Use/Cover Area frame Survey (LUCAS) (Ballin, Barcaroli, Palmieri and Zachariadis, 2025). Our method can be applied within strata to optimize rotation plans for a specific survey cycle. However, since auxiliary data and stratification often evolve over time, our approach is best viewed as a tool to optimize a fixed planning horizon, rather than for indefinite sample management.

Finally, the framework opens several possibilities for future algorithmic improvement. The targeted positional swap implemented in this work is an effective general-purpose search strategy, as it correctly focuses the search on the regions of the sequence that perform poorly. Although this type of swap is better than a blind swap, its random nature within the targeted region means that it may still have a low probability of suggesting an improvement. As a result, the number of iterations needed to find a really good sequence may be large.

The next significant step forward would be to create objective-specific, diagnostic-driven, move proposals that are informed by the cause of the poor score. For spatial balance, a high score for a panel indicates a large variance in its polytope masses, reflecting both overly dense and sparse regions in the auxiliary space. A truly intelligent move would therefore be a swap designed to directly remedy this flaw: removing a unit from one of the dense regions and replacing it with an outside unit that can better cover a sparse region. Similarly, for marginal balance, a diagnostic approach would identify the units that should be removed and the type of units that need to be added to improve the balance.

These diagnostic-driven heuristics, which propose moves with a high probability of being beneficial, represent a more sophisticated class of search and constitute a rich area for future research.

As this framework can be used to select master samples in which the panels meet different design criteria, such as spatial balance, marginal balance, or any other criteria, the estimation of variance using the resulting panels has not yet been investigated. It is likely that commonly used approximate variance estimators for spatially balanced samples would also work well for spatially balanced panels from a master sample. Similarly, for marginally balanced panels, an approximate variance estimator for balanced sampling is likely to work well. However, it is left for future investigation to determine exactly how the use of approximate variance estimators would perform for the different types of panels.

To facilitate the application of this framework, the R code used for the demonstrations in Section 4 is available on GitHub (Grafström, 2026). Given the modularity of the objective function, which can be tailored to any quantifiable design criterion, this code is intended to serve as an adaptable template. This approach provides users with flexibility to implement and optimize solutions for their specific needs.

References

Ballin, M., Barcaroli, G., Palmieri, A. and Zachariadis, S. (2025). Introduction of a panel component in the LUCAS survey and the use of a composed estimator for statistics production – 2025 edition. *Statistical Working Papers*, Publications Office of the European Union, Luxembourg.

Deville, J.-C. and Tillé, Y. (2004). Efficient balanced sampling: The cube method. *Biometrika*, 91(4), 893-912.

Eustache, E., Jauslin, R. and Tillé, Y. (2022). Spatiotemporal sampling with spatial spreading and rotation of units in time. *Spatial Statistics*, 47, 100613.

Grafström, A. (2012). Spatially correlated Poisson sampling. *Journal of Statistical Planning and Inference*, 142(1), 139-147.

Grafström, A. (2026). Master Samples with Optimized Panels [Computer software]. GitHub. Available at: https://github.com/antongrafstrom/MasterSamples.

Grafström, A., Lundström, N.L. and Schelin, L. (2012). Spatially balanced sampling through the pivotal method. *Biometrics*, 68(2), 514-520.

Grafström, A. and Prentius, W. (2024). `BalancedSampling`*: Balanced and Spatially Balanced Sampling*. R package version 2.1.1.

Grafström, A. and Tillé, Y. (2013). Doubly balanced spatial sampling with spreading and restitution of auxiliary totals. *Environmetrics*, 24(2), 120-131.

Grafström, A., Zhao, X., Nylander, M. and Petersson, H. (2017). A new sampling strategy for forest inventories applied to the temporary clusters of the Swedish national forest inventory. *Canadian Journal of Forest Research*, 47(9), 1161-1167.

Kirkpatrick, S., Gelatt Jr., C.D. and Vecchi, M.P. (1983). Optimization by simulated annealing. *Science*, 220(4598), 671-680.

Lark, R.M. (2016). Multi-objective optimization of spatial sampling. *Spatial Statistics*, 18, 412-430.

Larsen, D.P., Olsen, A.R. and Stevens, D.L. (2008). Using a master sample to integrate stream monitoring programs. *Journal of Agricultural, Biological, and Environmental Statistics*, 13(3), 243-254.

O'Luing, M., Prestwich, S. and Tarim, S.A. (2022). A simulated annealing algorithm for joint stratification and sample allocation. *Survey Methodology*, 48(1), 225-249. Available at: https://www150.statcan.gc.ca/n1/en/pub/12-001-x/2022001/article/00010-eng.pdf.

Prentius, W. (2024). Locally correlated Poisson sampling. *Environmetrics*, 35(2), e2832.

Prentius, W. and Grafström, A. (2024). How to find the best sampling design: A new measure of spatial balance. *Environmetrics*, 35(7), e2878.

Robertson, B., Price, C. and Reale, M. (2024). Well-spread samples with dynamic sample sizes. *Biometrics*, 80(2), ujae026.

Robertson, B., Price, C., Reale, M. and Davies, P. (2026). Doubly balanced samples with dynamic sample sizes. *Biometrics*, 82(1), ujag011.

Stevens Jr., D.L. and Olsen, A.R. (2004). Spatially balanced sampling of natural resources. *Journal of the American statistical Association*, 99(465), 262-278.

Tillé, Y., Dickson, M.M., Espa, G. and Giuliani, D. (2018). Measuring the spatial balance of a sample: A new measure based on Moran's I index. *Spatial Statistics*, 23, 182-192.

van Dam-Bates, P., Gansell, O. and Robertson, B. (2018). Using balanced acceptance sampling as a master sample for environmental surveys. *Methods in Ecology and Evolution*, 9(7), 1718-1726.