

OPTIMAL EXPERIMENTAL DESIGN OF MEASUREMENT CONFIGURATIONS FOR GEODETIC METROLOGY

M. Del Grande^{*1}, J. Gutekunst¹, and B. Yang¹

¹CERN, Geneva, Switzerland

Abstract

In high-precision metrology and geodetic network design, the achievable accuracy depends not only on the instruments but also on the measurement geometry – in particular, on where fiducials and instrument stations are placed. The choice of these positions is therefore a key lever for minimising the measurement uncertainty, yet it must be fixed before any measurement is taken. This contribution presents a framework, built on CERN’s geodetic least-squares adjustment platform LGC (Logiciel Général de Compensation), that applies Optimum Experimental Design (OED) to this problem. As the covariance of a least-squares estimate depends only on the Jacobian of the observation model at the nominal configuration and on the stochastic model, the quality of a candidate design can be graded from a single Jacobian evaluation, without any measurement and without a nonlinear adjustment. The estimation variables are split into quantities of interest, which define the objective function – D-optimality (error-ellipsoid volume) or A-optimality (average parameter variance) – and bounded design variables that the optimiser searches over. Two problem families are treated: a continuous one, in which station and fiducial coordinates are optimised over a bounded domain, and a discrete one, in which the most informative subset of observation groups is selected under a budget. The framework draws on LGC’s modelling capabilities, in particular its frames (local coordinate systems) parametrised by estimable parameters. We illustrate it on a fiducialisation scenario, optimising the placement of fiducials on magnet components and the positioning of the stations to determine the components’ relative position and orientation as precisely as possible, demonstrating its use as a design-phase support tool.

INTRODUCTION

The Geodetic Metrology (GM) group of CERN’s Beams Department is responsible for the metrology and alignment of the components installed in the accelerators, in their transfer lines, and in the physics experiments [1].

Achieving the stringent alignment specifications required in these environments heavily depends on the geometric configuration of the measurement networks. A primary example of this is the *fiducialisation* process, the laboratory operation that links a component’s functional axis to a set of external reference targets (fiducials). The outcome is an estimated Helmert frame. As fiducialisation errors cannot be observed or corrected in the tunnel, they propagate integrally

into the final installed position. Controlling the precision of this estimated frame, governed by the placement of fiducials and the configuration of instrument positions, is therefore critical.

Currently, these design choices are largely dictated by mechanical constraints, accessibility, and accumulated experience. Configurations are typically validated *a posteriori*: a candidate setup is simulated, its covariance matrix is inspected, and the design is either accepted or manually adjusted. While sound, this trial-and-error procedure explores a limited fraction of the configuration space. It can be a very time consuming approach, especially if the demanded accuracy is at the limit of what is technically feasible with the available instrumentation.

The starting point for a more systematic treatment is a classical property of least-squares estimation [2]: the covariance of the estimated parameters depends only on the Jacobian A of the observation model evaluated at the nominal configuration and on the stochastic model P , entirely independent of the actual observed values. The information content of a design is thus determined purely by its geometry and instrumentation. It can be rigorously evaluated before any physical measurement is taken at the cost of a single Jacobian evaluation rather than a full nonlinear adjustment. This is the foundation of Optimal Experimental Design (OED), formalized in the geodetic community as the Grafarend network optimization hierarchy [3] and in statistics as the theory of optimal designs [4, 5].

Crucially, the potential of this OED approach extends far beyond fiducialisation. It provides a universal, quantitative framework for any metrological scenario that can be formulated using bounded design and optimization variables. Highly impactful applications include general network planning to strictly control uncertainty propagation [6], as well as advanced R&D projects, such as optimizing sensor distributions and geometries for Frequency Scanning Interferometry (FSI) systems.

Moving beyond the validation of individual candidate configurations, this contribution applies OED systematically to the GM group’s design phase. We present a novel, object-oriented Python package built directly on the API interface of LGC [7], CERN’s primary software for most survey and alignment tasks. By leveraging the robust mathematical modelling capabilities and the analytical engine of LGC, this library translates theoretical OED principles into a highly impactful, practical design tool, empowering surveyors to mathematically optimize and guarantee their measurement configurations prior to any physical deployment.

^{*} mattia.delgrande@cern.ch

MATHEMATICAL FORMULATION

Gauss–Markov Model

Let $x \in \mathbb{R}^{n_x}$ denote the unknown parameters and $\ell \in \mathbb{R}^{n_\ell}$ the observations. The functional and stochastic models are defined as [2]:

$$\ell = f(\tilde{x}) + e, \quad e \sim \mathcal{N}(0, \sigma_0^2 P^{-1}) \quad (1)$$

where f is a differentiable observation model, P the weight matrix, and σ_0^2 the a priori variance factor. This formulation is highly versatile: polar observations, image rays, or interferometric distances all enter the same system.

The estimate is the solution of the nonlinear weighted least-squares problem

$$\hat{x} := \arg \min_{x \in \mathbb{R}^{n_x}} \|\ell - f(x)\|_P^2, \quad \|v\|_P^2 := v^\top P v, \quad (2)$$

which LGC solves via Gauss–Newton iterations. Linearising f at the solution $\hat{x}$ and writing

$$A = \frac{\partial f}{\partial x}(\hat{x}), \quad N = A^\top P A \quad (3)$$

Assuming full rank, first-order error propagation yields the covariance matrix $\Sigma_x = \sigma_0^2 N^{-1}$.

As (3) depends exclusively on the nominal geometry and a priori precisions, this formulation allows the information content of any candidate design to be rigorously evaluated before taking a single measurement

Quantities of interest and the concentrated information matrix

In design applications one is rarely interested in the full parameter vector. The quality of a configuration is judged on a distinguished subvector of unknowns—the *quantities of interest* (QoI)—while the remaining unknowns are carried only because the model requires them.

Assume throughout this subsection that the model is well posed. After a suitable reordering of the components of x , we obtain:

$$x = \begin{pmatrix} \theta \\ v \end{pmatrix}, \quad \theta \in \mathbb{R}^p, \quad v \in \mathbb{R}^q, \quad p + q = n_x, \quad (4)$$

where θ collects the quantities of interest and v the *nuisance parameters*. The partition is application-specific: θ may consist of the pose parameters ($T_X, T_Y, T_Z, R_X, R_Y, R_Z$) of a component frame—the leading case in this paper—but equally of the coordinates of selected points. A QoI defined as a linear combination $L^\top x$ reduces to the form (4) after a linear reparameterisation, so no generality is lost. Conformally with (5), split the Jacobian columnwise and the normal matrix blockwise [8, 9],

$$A = \begin{pmatrix} A_\theta & A_v \end{pmatrix}, \quad N = \begin{pmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{pmatrix} = \begin{pmatrix} A_\theta^\top P A_\theta & A_\theta^\top P A_v \\ A_v^\top P A_\theta & A_v^\top P A_v \end{pmatrix} \quad (5)$$

The *QoI covariance* is the leading $p \times p$ block of the full parameter covariance,

$$\Sigma_\theta := \sigma_0^2 [N^{-1}]_{11}. \quad (6)$$




Figure 1: Block partition of the scaled inverse normal matrix.

Optimality criteria

An *optimality criterion* is a scalar function Φ , to be minimised, that maps the symmetric positive-definite covariance matrix $\Sigma_\theta(w)$ of the quantities of interest into a scalar value. Geometrically, its natural support is the error ellipsoid of the estimator. Each of the three classical criteria [4, 5] controls a different physical feature of this error ellipsoid, as illustrated in Figure 2.

D-optimality.

$$\Phi_D(\Sigma_\theta) := \det \Sigma_\theta \quad (7)$$

Minimises the volume of the error ellipsoid, thereby reducing the overall uncertainty of the quantities of interest.

A-optimality.

$$\Phi_A(\Sigma_\theta) := \text{tr} \Sigma_\theta = \sum_{i=1}^p \lambda_i(\Sigma_\theta) \quad (8)$$

Minimises the sum of the variances (the trace of the matrix), which corresponds to the expected squared Euclidean norm of the estimation error, or the mean square extent of the ellipsoid.

E-optimality.

$$\Phi_E(\Sigma_\theta) := \lambda_{\max}(\Sigma_\theta) \quad (9)$$

Minimises the variance in the worst-determined direction, effectively shrinking the longest semi-axis of the ellipsoid.




Figure 2: Geometric interpretation of the three classical optimality criteria on the error ellipsoid.

CONTINUOUS FORMULATION

Having established the scalar criteria that evaluate the information content of a given configuration, the focus now shifts to optimising them. In the continuous formulation the observation plan is fixed: which instrument observes which target, of which type, while the coordinates at which the model is linearised are free. This section formalises the mapping from this continuous geometric space to the resulting estimation precision, casts the search for the optimal layout as a bound-constrained minimisation problem, and outlines the three numerical strategies used to solve it.

The design map

Let $\mathbf{w} \in \mathcal{S} \subseteq \mathbb{R}^m$ denote the vector of *design variables*: the free geometric quantities of the measurement configuration, such as the positions of the instrument stations and of the fiducial targets.

The precise structure of $\mathbf{w}$ and of the feasible set $\mathcal{S}$ depends on the formulation. Mathematically, $\mathbf{w}$ corresponds to a specific subset of the estimation variables — typically drawn from the nuisance parameters $\mathbf{v}$ — meaning it effectively lives in a subspace of the full parameter vector $\mathbf{x}$. For the present section, it suffices that the nominal geometry, and hence the linearised observation model of Equation (1), is a function of $\mathbf{w}$.

Every design induces the composition

$$\mathbf{w} \mapsto A(\mathbf{w}) \mapsto N(\mathbf{w}) \mapsto \Sigma(\mathbf{w}) \mapsto \Sigma_\theta(\mathbf{w}) \quad (10)$$

which we call the *design map*. Its first arrow is the only non-linear and problem-specific one: it evaluates the Jacobian with respect to the estimation variables at the nominal configuration dictated by $\mathbf{w}$. The subsequent arrows are fixed algebraic operations and structural extractions.

Formulation

The design map 10 carries a design $\mathbf{w}$ to the full parameter covariance $\Sigma(\mathbf{w})$. Composing it with the extraction of the QoI and with a criterion of the previous section yields a scalar objective on the design space,

$$f(\mathbf{w}) := \Phi(\Sigma_\theta(\mathbf{w})) \quad \Phi \in \{\Phi_D, \Phi_A, \Phi_E\}. \quad (11)$$

The continuous design problem is the minimization of Eq. (11) over the configurations that the mechanical environment allows.

Let the feasible set be the box

$$\mathcal{S} = \prod_{j=1}^m [l_j, u_j] \subset \mathbb{R}^m, \quad l_j < u_j, \quad (12)$$

whose bounds encode the reachable range of each design coordinate. The continuous optimal design problem is

$$\min_{\mathbf{w} \in \mathcal{S}} f(\mathbf{w}) = \min_{\mathbf{w} \in \mathcal{S}} \Phi(\Sigma_\theta(\mathbf{w})). \quad (13)$$

We solve this problem with numerical algorithms, such as L-BFGS [10], but we cannot guarantee to reach the global optimum because of non-convexity [11].

DISCRETE FORMULATION

In contrast to the continuous problem, the discrete setting features a fully fixed geometry and a varying set of candidate observations. Because we only select from this predefined set, a single Jacobian evaluation suffices in this process. Freezing the spatial coordinates effectively bypasses the need for repeated, computationally expensive Jacobian evaluations. The problem then shifts entirely to determining *which* observations to carry out. Constrained by operational measurement budgets, the goal is to select the most informative subset of observations from a complete catalog (e.g. instrument set-ups, photogrammetric frames) without exceeding a given quota.

Formulation

Group structure and additivity. Candidate observations often cannot be selected individually because they are operationally inseparable (e.g., a photogrammetric frame delivers all its image rays at once, and an instrument set-up yields a full round of measurements). Therefore, the observation catalogue is partitioned into G disjoint candidate groups, alongside a fixed base set $\mathcal{B}_0$ that is always measured:

$$\{1, \dots, n_\ell\} = \mathcal{B}_0 \sqcup \mathcal{G}_1 \sqcup \dots \sqcup \mathcal{G}_G. \quad (14)$$

Because the observations are uncorrelated, the information matrices accumulate additively. The normal matrix for any selected subset $\mathcal{S}$ of candidate groups is simply:

$$N(\mathcal{S}) = N_0 + \sum_{g \in \mathcal{S}} N_g \quad (15)$$

where $N_i = A_i^T P A_i$ for $i \in \{0, 1, \dots, G\}$.

This additivity is the structural core of the discrete formulation: a candidate group is graded purely by the independent information it adds to the accumulated system, allowing for rapid and incremental evaluations.

Objective. The covariance of the quantities of interest, $\Sigma_\theta(\mathcal{S})$, is extracted from the inverted normal matrix $N(\mathcal{S})^{-1}$. Given a budget K of affordable observation groups drawn from a candidate pool Ω , the goal is to find the measurement subset that yields the highest precision. Provided the resulting network geometry is strictly well-posed, the discrete optimal design problem is directly cast as the constrained minimisation of the chosen optimality criterion:

$$\min_{\mathcal{S} \subseteq \Omega, |\mathcal{S}|=K} \Phi(\Sigma_\theta(\mathcal{S})) \quad \text{subj to } N(\mathcal{S}) \succ 0. \quad (16)$$

Consequently, we deploy greedy algorithms [12] to traverse the search space. While these heuristic methods do not guarantee strict global optimality, they efficiently identify highly informative measurement subsets in a fraction of the time, bridging the gap between theoretical optimality and practical surveying needs.

COMPUTATIONAL STRATEGIES

Evaluating the optimality criteria naively is computationally prohibitive for realistically sized networks, as it requires extracting covariances across thousands of trial configurations. To overcome this bottleneck, we introduce three algorithmic strategies: the first applies broadly to both formulations, the second is tailored to continuous placement, and the third accelerates discrete selection.

Eliminating the Nuisance Parameters

A geodetic adjustment estimates far more parameters than the quantities of interest θ (e.g., the pose of a specific magnet or small assembly). The normal matrix also carries instrument positions, orientation unknowns, and auxiliary network points. While these nuisance parameters cannot be ignored—since they correlate with θ and omitting them leads to an overly optimistic uncertainty—their individual values are ultimately irrelevant to the design optimization.

Partitioning the parameter vector splits the normal matrix into:

$$N = \begin{bmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{bmatrix}, \quad (17)$$

where $N_{11} \in \mathbb{R}^{k \times k}$ and $N_{22} \in \mathbb{R}^{(n-k) \times (n-k)}$. Computing the desired covariance naively requires inverting the full $n \times n$ matrix—a costly $\mathcal{O}(n^3)$ operation—just to extract the top-left $k \times k$ block, where $k \ll n$.

The Schur complement avoids this overhead by computing the reduced matrix directly [8, 9]:

$$S = N_{11} - N_{12} N_{22}^{-1} N_{21} \quad (18)$$

By the properties of block matrix inversion, this establishes the explicit identity $S^{-1} = [N^{-1}]_{11}$. Matrix S therefore represents the true information surviving on θ after optimally accounting for the nuisances.

From a computational standpoint, the inverse matrix N_{22}^{-1} is never explicitly formed. Instead, the reduction only requires solving linear systems for the k columns of N_{21} , each of dimension $n - k$, which is significantly less work. Exploiting the inherent sparsity of N_{22} via Cholesky decomposition keeps these column solves highly efficient [9].

Grouped Finite Differences by Graph Colouring

Our continuous formulation relies on gradient-based optimization. For the gradient of the objective, we need the sensitivity of certain entries of the Jacobian via finite differences. Naively, evaluating a design vector $x \in \mathbb{R}^d$ demands $d + 1$ full network evaluations per step, creating a severe computational bottleneck.

However, this process is highly redundant. A design variable only affects the specific observations it physically involves. Therefore, variables that never share an observation are mutually independent and can be perturbed simultaneously without cross-contamination. For an independent set




Figure 3: Conflict graph of a small example network; an edge joins two design variables that share an observation. Four colours suffice for a conflict-free colouring, so one gradient step costs $4+1$ instead of $12+1$ network evaluations.

of variables C , the combined perturbation is:

$$\Delta J = J \left(x + h \sum_{i \in C} e_i \right) - J(x) \quad (19)$$

where e_i are the standard basis vectors. The individual sensitivities $\partial J_{m,i} / \partial x_i$ can then be cleanly extracted from their respective rows.

Finding a partition of the variables into independent, conflict-free subsets can be formulated as a graph colouring problem [13, 14], which is solved efficiently by standard libraries. Variables assigned the same colour are independent and perturbed together. Since the network topology is fixed, this colouring algorithm is executed only once at initialization. For the sparsity patterns arising in geodetic networks, this drastically reduces the number of evaluations needed for the finite-difference gradient—from roughly one per design variable naively, to fewer than 10 even for thousands of points.

Accelerated Candidate Selection via Rank- r Updates

In the discrete selection problem, the greedy algorithm evaluates hundreds of candidate observation groups per iteration. Naively reconstructing and inverting the $k \times k$ information matrix for each candidate scales cubically, creating a severe computational bottleneck for realistically sized networks.

To overcome this, we exploit the fact that adding a group of r observations ($r \ll k$) is fundamentally a *rank- r update*. We leverage the Matrix Determinant Lemma to rapidly score the information gain of each candidate, and the *Woodbury identity* to update the inverse matrix once the optimal candidate is chosen [8, 9].

This approach restricts matrix operations exclusively to the highly efficient $r \times r$ space of the new observations, bypassing the massive $k \times k$ parameter space, this strategy delivers orders-of-magnitude acceleration and allows the optimization to scale effortlessly to massive measurement configurations.

CASE STUDIES

In this section, we analyze how the proposed quantitative approach can assist surveyors in their practical tasks. As discussed in the introduction, the core process under examination is fiducialisation. Specifically, two distinct scenarios are investigated: the fiducialisation required for the alignment of two magnets (a quadrupole with 2 apertures and two sextupoles), and the measurement of reference targets on a girder using Photogrammetry technology.

Quadrupole and Sextupole Fiducialisation

In the first scenario, the objective is to assess the impact, of both the number and spatial distribution of the targets, as well as the positioning of the laser trackers used for the measurements, on the estimation of the frames.

This process is carried out sequentially. As illustrated in Figure 4, the assembly consists of three components: one quadrupole and two sextupoles. The ultimate goal is to align the two functional axes of the sextupoles with the respective ones of the quadrupole. This is achieved by estimating the local coordinate frames of the individual components through the measurement of targets located on their surfaces.




Figure 4: Representation of the Girder, Quadrupole, Sextupoles setup

For the geometric measurements in this scenario, a Leica Absolute Tracker AT930 was simulated. It should be noted that the AT930 is utilized here as a representative example; any instrument within the Leica AT family will yield comparable performance and results for this purpose. The measurements rely on 1.5-inch Break Resistant Reflectors. Within the LGC simulation environment, the 1-sigma standard deviations for this setup were defined as follows: 2.5 cc (1 cc = 10^{-4} gon) for the horizontal angle, 3 cc for the zenith distance, and 0.01 mm for the absolute distance measurement, coupled with a scale factor error of 0.5 ppm. All other adjustable variables, including distance corrections, target centering errors, and target height uncertainties, were assumed to be negligible and set to zero.

The procedure begins with the quadrupole: the optimization algorithm is executed to determine the optimal placement of the fiducials for the best estimation of the Helmert frame. Subsequently, a simulation is performed using the LGC software [7] to evaluate the associated uncertainty. This step determines the precision with which the quadrupole frame can be established.

The subsequent phase involves mounting the quadrupole onto the girder and repeating the procedure for the two sextupoles, ensuring the accurate alignment of all three frames. To accomplish this, the sextupole is defined with respect to the quadrupole frame, as the relative orientation between the frames is of primary interest, rather than their absolute positions in space. The algorithm is then executed again to compute the optimal positions of both the fiducials on the sextupole and the instrument stations in space, subject to the physical constraints and simple bounds on their allowed positions.




Figure 5: Representation of the Quadrupole-Sextupoles procedure

The aim is to investigate the impact of fiducial placement, considering both the quantity and the feasible installation locations. Specifically, the analysis focuses on how the in-

formation regarding the coordinate frames improves as the number of fiducials increases.

For the quadrupole, the following three configurations were analyzed: 4 fiducials on the top face; 4 fiducials on the top face and 4 on the lateral faces; and finally, 4 fiducials on the top face combined with 8 on the lateral faces. The corresponding results are as follows:

Table 1: Estimated uncertainties for the Quadrupole frame parameters across the three analyzed configurations.

Parameter	Unit	4 fiducials	8 fiducials	12 fiducials
TX	(mm)	0.0147	0.0080	0.0067
TY	(mm)	0.0109	0.0080	0.0069
TZ	(mm)	0.0110	0.0085	0.0073
RX	(cc)	4.691	3.561	2.766
RY	(cc)	23.941	12.605	9.615
RZ	(cc)	4.408	3.372	2.679

A similar approach was applied to the sextupole, evaluating the following configurations: 3 fiducials on the top face; 2 on the top face and 2 on the side faces; and 2 on the top face, 2 on the side faces, and 2 on the bottom face, as illustrated in Figure 6.




Figure 6: Highlighted faces where we can put the fiducials

The obtained results are presented below:

Table 2: Estimated uncertainties for the Sextupole frame parameters across different fiducial configurations.

Parameter	Unit	3 fiducials	4 fiducials	6 fiducials
TX	(mm)	0.0309	0.0098	0.0085
TY	(mm)	0.0090	0.0072	0.0062
TZ	(mm)	0.0106	0.0094	0.0090
RX	(cc)	19.685	13.759	11.039
RY	(cc)	109.754	27.868	21.056
RZ	(cc)	16.889	13.329	11.529

As demonstrated by the results in Table 1 and Table 2, the OED framework [3, 4] provides a rigorous, quantitative basis for evaluating the trade-off between measurement uncertainty and setup complexity. It is important to emphasize that these values are not merely the result of arbitrary simulations, but are the direct output of the optimization software, which computed the mathematically optimal instrument station placements for each specific fiducial configuration.

Analyzing the data, the most impactful finding is the behavior of the rotational parameters, particularly R_Y . In accelerator metrology, resolving R_Y is notoriously critical for the accurate alignment of the magnetic axes [1].




Figure 7: The final optimized configuration, demonstrating the strategic placement of fiducials and stations

Measurement with Photogrammetry

The objective of this scenario is to perform the measurement of a girder with the same setting as before (quadrupole and 2 sextupoles), utilizing photogrammetry. The primary goal is to determine the position of the targets with the lowest possible uncertainty, thereby minimizing the associated uncertainty to achieve accuracy levels comparable to those obtained with laser trackers.

To address this problem, an initial configuration comprising 722 well-distributed camera positions around an object covered with 108 targets was considered. To illustrate the spatial distribution and the sheer density of this baseline configuration, Figure 8 depicts the complete initial setup.




Figure 8: Visualization of the complete initial photogrammetric setup, comprising 722 potential camera positions distributed around the target object equipped with 108 targets.

Naturally, deploying such a high number of camera positions in a real-world scenario is practically impracticable due to the enormous volume of measurements required. This limitation directly motivates the application of our second proposed method: the discrete optimization approach. This technique allows for the selection of the k most informative camera positions, specifically tailored to our chosen Quantity of Interest (QoI).




Figure 9: Information gain as a function of the number of selected camera positions.

Figure 9 illustrates the sequential addition of cameras by the algorithm, up to the full set of 722. As can be observed, adding more cameras beyond approximately 400 yields a negligible gain in information, indicating a clear plateau. However, it is important to underline that while the information gain curve provides a valuable global overview of where the network saturates, what is ultimately of practical concern to surveyors are the actual coordinate uncertainties (the σ values) of the measured points. Furthermore, this first plot is derived from a metric that inherently carries information regarding the mathematical correlations between the variables, offering a holistic view of the network's quality.

To translate this optimization into tangible metrological terms, it is crucial to analyze the standard deviations along the three spatial components (S_X , S_Y , S_Z), which represent the respective measurement uncertainties along the coordinate axes. Figure 10 illustrates the behavior of the *average* uncertainty computed across all targets for different subsets of optimized cameras. While this second plot does not capture the correlations between variables—unlike the information gain curve—it provides a direct, interpretable metric of the expected geometric precision.

To visualize the physical distribution resulting from this approach, Figure 11 displays a strictly optimized network restricted to a budget of exactly 50 camera positions. This specific number was chosen because, as shown by the average standard deviations, the resulting measurement uncer-

tainty is already sufficiently low to meet our strict alignment tolerances. Despite the severe reduction from the initial 722 candidates, the algorithm strategically distributes these 50 selected camera positions to maintain a strong, intersecting network geometry across all targets.




Figure 11: Optimized photogrammetric network comprising a subset of 50 strategically selected camera positions. The network geometry maintains robust intersections despite the drastically reduced camera positions.

Ultimately, this approach enables an *a priori* analysis of various measurement scenarios. It serves as a highly valuable decision-support tool for surveyors, who can adjust the initial settings based on their practical experience. Simultaneously, it provides deeper insights into geometric and metrological aspects that might not be intuitively obvious, but are rigorously captured and evaluated by the mathematical model.

A possible future step could be to fix this optimal camera configuration and apply the same methodology to progressively reduce the number of photogrammetric targets on the objects, further streamlining the measurement process.

CONCLUSION

This contribution has presented an Optimal Experimental Design (OED) framework for geodetic measurement configurations, built on the Python interface of LGC2. Historically, the placement of fiducials and instrument stations has been largely dictated by mechanical constraints, accessibility, and accumulated experience. Previous validation procedures




Figure 10: Comparison of the average standard deviations (S_X , S_Y , S_Z) across all targets as a function of different optimized camera configurations.

explored only a limited fraction of possible configurations and offered no statement on how close a chosen setup was to the theoretical optimum. This novel quantitative approach bridges that gap, providing a layer of analysis previously missing in the Geodetic Metrology group.

By leveraging the property that the covariance of a least-squares estimate depends only on the Jacobian of the observation model at the nominal configuration and the stochastic model, a candidate design can now be evaluated exactly before any physical measurement is taken. Grading is performed on the marginal covariance of the quantities of interest, seamlessly accounting for nuisance parameters.

For the surveyor, this translates into a highly practical decision-support tool. The case studies demonstrate how the framework definitively answers geometric questions that were previously settled by intuition or trial and error alone. In magnet fiducialisation, the algorithm quantitatively demonstrated that the marginal gain of adding fiducials eventually saturates; intermediate configurations already achieve the necessary precision, and introducing lateral lever arms improves the dominant rotational component R_Y by a factor of two. Similarly, in the photogrammetric scenario, the tool revealed that a much lower number out of 722 candidate camera positions are needed to capture essentially all available information regarding the fiducial coordinates.

Ultimately, this framework acts as a pioneering design-phase asset. It does not replace the surveyor’s judgment, but rather empowers it by rigorously quantifying the trade-off between metrological performance and operational effort. While the current optimization focuses exclusively on measurement uncertainty, reliability remains a critical parameter for all network setups that should not be ignored. Consequently, extending the objective function to explicitly account for reliability is a key avenue for future research. Alongside this, ongoing work addresses visibility and accessibility constraints inside the feasible set, paving the way for the integration of this package into the routine fiducialisation workflow.

REFERENCES

- [1] J.-C. Gayde, “Survey and alignment of accelerators”, in *CERN Yellow Reports: Monographs*, 2024.
- [2] K.-R. Koch, *Parameter Estimation and Hypothesis Testing in Linear Models*, 2nd ed. Springer, 1999.
- [3] E. W. Grafarend, “Optimization of geodetic networks”, *Boll. Geod. Sci. Affini*, vol. 33, no. 4, 1974.
- [4] F. Pukelsheim, *Optimal Design of Experiments*. SIAM, 2006.
- [5] V. V. Fedorov, *Theory of Optimal Experiments*. Academic Press, 1972.
- [6] V. Gerligand, D. Mergelkuhl, L. Watrelot, and S. Durand, “Evaluation of error propagation in the first-order geodetic network of the FCC tunnel using topographical techniques”, in *Proc. IPAC’26*, Deauville, France, 2026, 2026, paper MOP7132, doi:10.18429/JACoW-IPAC2026-MOP7132
- [7] G. Kautzmann, J. Gutekunst, F. Klumb, R. Singh, and P. Elson, “LGC: CERN’s geodetic network adjustment software goes open source”, in *Proc. FIG Congress 2026*, Cape Town, South Africa, 2026, 2026, https://fig.net/resources/proceedings/fig_proceedings/fig2026/papers/ts08d/TS08D_kautzmann_gutekunst_et_al_13858.pdf
- [8] D. A. Harville, *Matrix Algebra From a Statistician’s Perspective*. Springer, 1997.
- [9] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed. Johns Hopkins University Press, 2013.
- [10] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. Springer, 2006.
- [11] C. G. E. Boender and A. H. G. R. Kan, “Bayesian stopping rules for multistart global optimization methods”, *Math. Program.*, vol. 37, no. 1, pp. 59–80, 1987.
- [12] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. MIT Press, 2009.
- [13] A. R. Curtis, M. J. D. Powell, and J. K. Reid, “On the estimation of sparse Jacobian matrices”, *J. Inst. Math. Appl.*, vol. 13, no. 1, pp. 117–119, 1974.
- [14] T. F. Coleman and J. J. Moré, “Estimation of sparse Jacobian matrices and graph coloring problems”, *SIAM J. Numer. Anal.*, vol. 20, no. 1, pp. 187–209, 1983.