

Towards Gauge Invariant Bundle Adjustment: A Solution Based on Gauge Dependent Damping

Adrien Bartoli

INRIA Rhône-Alpes, 655, avenue de l'Europe
38334 Saint Ismier cedex, France. *First.Last@inria.fr*

Abstract

Bundle adjustment is used to obtain accurate visual reconstructions by minimizing the reprojection error. The coordinate frame ambiguity, or more generally the gauge freedoms, has been dealt with in different manners. It has often been reported that standard bundle adjustment algorithms were not gauge invariant: two iterations within different gauges can lead to geometrically very different results. Surprisingly, most algorithms do not exploit gauge freedoms to improve performances. We consider this issue. We analyze theoretically the impact of the gauge on standard algorithms. We show that a sufficiently general damping matrix in Levenberg-Marquardt iteration can be used to implicitly reproduce a gauge transformation. We show that if the damping matrix is chosen such that the decrease in the reprojection error is maximized, then the iteration is gauge invariant. Experimental results on simulated and real data show that our gauge invariant bundle adjustment algorithm outperforms existing ones in terms of stability.

1. Introduction

The recovery of accurate 3D structure and camera motion from images is a major research challenge in photogrammetry [1] and computer vision [17]. Bundle adjustment is a technique to refine a visual reconstruction to produce jointly optimal 3D structure and camera motion. Under certain assumptions on the noise on the observed features, bundle adjustment consists in minimizing the reprojection error, which is in general a non-linear procedure. Second-order non-linear least squares algorithms are usually employed, namely the Gauss-Newton and Levenberg-Marquardt methods. These methods iteratively improve sub-optimal parameter estimates by solving *normal equations*. Efficient solutions are possible thanks to the sparse block structure of the normal equations. The Levenberg-Marquardt method has proved the most successful due to the use of a trust region strategy, implemented via a *damping* of the normal equations, also called *augmentation*.

An inherent problem in bundle adjustment is the choice of the coordinate frame in which the reconstruction is expressed. This coordinate frame is called the *gauge*, or the datum in the photogrammetry community. A gauge is a subset of parameter vectors such that any two of them do not share the same underlying geometry. Parameter vectors corresponding to the same geometry are related by *gauge transformations*. It has been shown in [8] that these transformations have a group structure. The reprojection error is gauge invariant since it reflects the inherent merit of a reconstruction. These considerations hold whichever camera model is used and whichever calibration level is available. Examples of *geometric gauge freedoms* are the global rotation, translation and scale in metric reconstruction and the position of points along the line in the two point representation of lines. Example of *algebraic gauge freedoms* is the scale factor of homogeneous coordinates. When ignored, gauge freedoms may induce numerical estimation problems, since they imply the rank deficiency of the normal equations.

It has been reported that most bundle adjustment algorithms initialized with the same reconstruction expressed within different gauges can yield extremely different results, in terms of computational cost [1, 3, 8, 11, 12, 14, 17]. In other words, these algorithms are *not gauge invariant*¹. Carefully choosing the gauge can greatly improve the reliability. In practice, we observe that most algorithms converge to the same solution, but may require different number of iterations.

Gauge freedoms imply that the normal equations to be solved at each iteration do not have a unique solution, i.e. the design matrix is singular, but rather a family of solutions, whose dimension is the number of gauge freedoms. Damping the normal equations consists in adding a symmetric positive definite matrix to the design matrix. This matrix is most

¹We employ 'gauge invariant' in a manner different from [8, 11, 12]. These authors propose algorithms based on using a pre-defined global gauge. Hence, no matter the gauge within which the initial solution is expressed, their algorithms will give the same result. However, the behaviour of these algorithms do depend upon the pre-defined global gauge. In this sense, these algorithms are *not* gauge invariant.

of the time chosen as a diagonal one, which corresponds to an elliptical trust region, see e.g. [16] and §3 for more details.

Surprisingly, gauge freedoms have rarely been exploited to improve on the reliability of algorithms. Most algorithms are not gauge invariant and are based on somehow arbitrary gauges which are not chosen to meet some efficiency criteria. More details are given in §2.

The main contribution of this paper is an algorithm that maximizes the decrease in the reprojection error at each iteration, in a gauge invariant manner. In more detail:

- In §4, we formally state the gauge invariance property of a bundle adjustment algorithm. A detailed investigation of how the gauge influences Gauss-Newton and Levenberg-Marquardt iterations allows to conclude that in general neither of them is gauge invariant. This analysis shows that the gauge and the damping matrix in Levenberg-Marquardt iterations are closely linked: Gauge transformations change the standard elliptical trust region to a more complex shape.
- In §5, we concentrate on Levenberg-Marquardt iterations and propose *implicit gauges*, as opposed to the explicit transfer of a reconstruction to the desired gauge. This shows that in some sense, the damping matrix can encapsulate gauge transformations. We state the following important result. If the form of the damping matrix is sufficiently general, then the iteration can be made gauge invariant by carefully choosing this matrix, based on a gauge dependent criterion.
- In §6, we propose a gauge dependent damping matrix which defines a trust region with the following important properties. (i) it makes the underlying iteration gauge invariant, (ii) it maximizes the decrease in the reprojection error up to first order and (iii) it preserves the sparse block structure of the normal equations.

Finally §7 reports some experimental results and §8 concludes. For simplicity of notation, we deal with projective reconstruction of points and cameras. Extension to other camera models (e.g. affine cameras), other calibration levels (e.g. metric reconstruction) and other types of features (e.g. lines) is straightforward, following [8, 12]. Our final *gauge invariant bundle adjustment* algorithm is summarized in a practical manner in table 1.

Notation. We make no formal distinction between coordinate vectors and physical entities. Equality up to a non-null scale factor is denoted by \sim and is sometimes written by introducing explicitly the scale factor: $(\mathbf{x} \sim \mathbf{x}') \Leftrightarrow (\exists \alpha | \mathbf{x} = \alpha \mathbf{x}')$. Transposition and transposed inverse are denoted by \top and $^{-\top}$. Vectors are typeset using bold fonts (\mathbf{q} , \mathbf{Q}), matrices using sans-serif fonts (P , T) and scalars in italics. Indices are used to indicate the size of a matrix or vector ($P_{(3 \times 4)}$,

$\mathbf{q}_{(3 \times 1)}$) or to index a set of entities. The row-wise matrix vectorization is written vect .

Let $P_i, i = 1 \dots n$ denote the n reconstructed camera matrices and $\mathbf{Q}_j, j = 1 \dots m$ denote the m reconstructed 3D points. Structure and motion parameters are contained in a $(p \times 1)$ vector \mathcal{X} , partitionned into $12n$ motion parameters \mathcal{M} and $4m$ structure parameters \mathcal{S} (homogeneous coordinates of m points) as:

$$\mathcal{X}^\top = \underbrace{(\text{vect}^\top(P_1) \dots \text{vect}^\top(P_n))}_{\mathcal{M}^\top} \underbrace{(\mathbf{Q}_1^\top \dots \mathbf{Q}_m^\top)}_{\mathcal{S}^\top}$$

Changing the 15-degrees of freedom projective reconstruction basis and the individual scale factors of the camera matrices and point coordinate vectors leave the underlying geometry invariant. Hence there are $g = 15 + n + m$ degrees of gauge freedom.

2. Previous work

A survey of gauge freedom handling strategies is [17, §9]. Methods to deal with this problem are either *globally free*, i.e. they left the gauge free to drift, or *globally fixed*, i.e. they enforce a global gauge.

Globally free methods are based on selecting a solution of the normal equations, using a pseudo-inverse [1], numerical damping [9, 10, 16, 17] or local gauge constraints. For example, the standard photogrammetric inner constraints specify that the reconstructed points should not be translated, rotated and scaled, but they do not specify where are the reconstructed points, see [3, 17]. They are devised to minimize the norm of the parameter vector update, as well as the variance of the parameter estimate. In [1, §9.5], it is shown experimentally that this behaves better than globally fixed methods (see below), but that leaving the gauge drift freely can sometimes be better.

Globally fixed methods enforce a pre-chosen gauge. Some methods use reference elements (trivial gauge), either features [1, 4, 13] (object-centered gauge) or cameras [2, 5, 6]. In [1, p.40] the author suggests that control points or reference lengths can be used to fix a metric reconstruction basis, while in [4, 13] a similar method based on 5 reference points is used to fix a projective reconstruction basis. The constraints are enforced using either Lagrange multipliers [1, p.40] or by elimination [1, p.41], [4, 13]. In [6], the author considers projective reconstruction and partially fix the gauge by setting a canonical form to a reference camera, while [2] uses two reference cameras.

Another possibility for a globally fixed gauge is to use global gauge constraints [8, 11, 12], which raises the problem of enforcing these constraints during optimization. In [11], artificial extra observations are added to the cost function with a heavy weight. In [8], the gauge is left free to drift during the iteration, but the result is projected onto the gauge afterward, while in [12], a parameter subspace projected onto

the gauge is computed before the iteration. Another possibility is [1, p.41], where Lagrange multipliers are introduced and incorporated as unknowns in the optimization.

Besides the globally free photogrammetric inner constraints, all these methods are based on somehow arbitrary gauges which are not chosen to meet some efficiency criteria. A recent attempt in this direction is [14]. Based on gauge theory, the authors analyze, in the context of Euclidean reconstruction, which physical measure is the most likely to maximize the accuracy of the reconstruction while setting the scene scale.

3. Standard Algorithms

Bundle adjustment consists in refining a visual reconstruction to produce jointly optimal 3D structure and camera motion. By optimal, we mean that the maximum likelihood estimate of structure and motion is sought, which is achieved by minimizing an appropriate cost function. If we assume that the observed point features are corrupted by independently and identically distributed Gaussian noise, then the cost function is the reprojection error, given by the sum of squared differences between observed features \mathbf{q}_{ij} and predicted features $\hat{\mathbf{q}}_{ij}$:

$$\mathcal{C}(\mathcal{X}) = \sum_{i=1}^n \sum_{j=1}^m w_{ij} d^2(\mathbf{q}_{ij}, \hat{\mathbf{q}}_{ij}) = \mathbf{r}^T \mathbf{r}, \quad (1)$$

where w_{ij} is one if point j is visible in view i and zero otherwise. The $(r \times 1)$ vector \mathbf{r} is the residual error vector. Predicted features are given by $\hat{\mathbf{q}}_{ij} \sim \mathbf{P}_i \mathbf{Q}_j$.

Given a sub-optimal solution \mathcal{X}_0 , the parameter vector \mathcal{X} is iteratively updated by $\mathcal{X} \leftarrow \mathcal{X} + \delta$, where the increment δ is obtained as follows. Let $\mathbf{J} = \frac{\partial \mathbf{r}}{\partial \mathcal{X}}$ denotes the $(r \times p)$ Jacobian matrix of the residual error vector \mathbf{r} with respect to structure and motion parameters \mathcal{X} , $\mathbf{g} = \mathbf{J}^T \mathbf{r}$ be the $(p \times 1)$ gradient vector of \mathcal{C} and let $\mathbf{N} = \mathbf{J}^T \mathbf{J}$ be the Gauss-Newton approximation of the $(p \times p)$ Hessian matrix. The reprojection error is approximated by:

$$\mathcal{C}(\mathcal{X} + \delta) \approx \mathcal{C}(\mathcal{X}) + \underbrace{\mathbf{g}^T \delta + \frac{1}{2} \delta^T \mathbf{N} \delta}_e. \quad (2)$$

The minimum of this simple local quadratic model can be found by setting $\frac{\partial e}{\partial \delta} = 0$, which gives the Gauss-Newton iteration, through the following normal equations:

$$\mathbf{N} \delta = -\mathbf{g}. \quad (3)$$

Note that \mathbf{N} has a rank deficiency of g (the number of gauge freedoms). The Levenberg-Marquardt iteration is based on *damping* or augmenting the normal equations, as follows :

$$\underbrace{(\mathbf{N} + \mathbf{W}(\lambda))}_{\mathbf{N}^*} \delta = -\mathbf{g}, \quad (4)$$

where $\lambda > 0$ is related to the trust region radius and $\mathbf{W}(\lambda)$ is some symmetric positive definite $(p \times p)$ weight matrix, called the *damping matrix*, often chosen as:

$$\mathbf{W}(\lambda) = \lambda \mathbf{I}_{(p \times p)},$$

which corresponds to a spherical trust region. This is the original strategy proposed by Levenberg and Marquardt [9, 10] and recommended in [17]. Another commonly used solution, due to [16] and recommended in [5, 7, 15] is $\mathbf{W}(\lambda) = \lambda \mathbf{D}$ where \mathbf{D} is a diagonal matrix containing the diagonal entries of \mathbf{N} , which gives an elliptical trust region. The damping matrix must satisfy a normalization constraint so that the trust region radius is meaningful, e.g. $\|\mathbf{I}_{(p \times p)}\| = \sqrt{p}$. Note that the damping guarantees that \mathbf{N}^* is full-rank.

Parameter λ is tuned as follows: if parameters $\mathcal{X} + \delta$ decrease the error, i.e. if $\mathcal{C}(\mathcal{X} + \delta) < \mathcal{C}(\mathcal{X})$, then the step is accepted and the value of λ is divided by some constant, often 10, else the step is rejected and λ is multiplied by the constant.

4. Influence of the Gauge

Not surprisingly, It has been reported that in general, bundle adjustment was not gauge invariant, i.e. an iteration within different gauges yield geometrically different results [1, 3, 8, 11, 12, 14, 17]. We formulate the gauge invariance property and examine under which condition standard algorithms satisfy it.

4.1. Gauge Transformation and Invariance

Gauge transformations change the parameter vector without changing the underlying geometry, e.g. [8]. Let $\mathcal{X}^T = (\mathcal{M}^T \ \mathcal{S}^T)$ and $\check{\mathcal{X}}^T = (\check{\mathcal{M}}^T \ \check{\mathcal{S}}^T)$ be two parameter vectors of the same reconstruction expressed within two different gauges \mathcal{G} and $\check{\mathcal{G}}$. Let \mathbf{T} be the full-rank transformation relating the two underlying projective bases, defined such that: $\check{\mathbf{P}}_i = \gamma_i \mathbf{P}_i \mathbf{T}$ and $\check{\mathbf{Q}}_j = \alpha_j \mathbf{T}^{-1} \mathbf{Q}_j$, where γ_i and α_j are unknown non-zero scale factors. Entities written with a $\check{}$ are expressed within the gauge $\check{\mathcal{G}}$. The structure and motion parameters transform as:

$$\begin{aligned} \check{\mathcal{S}} &= \text{diag}(\underbrace{\alpha_1 \mathbf{T}^{-1}, \dots, \alpha_m \mathbf{T}^{-1}}_m) \mathcal{S} \\ \check{\mathcal{M}} &= \text{diag}(\underbrace{\gamma_1 \mathbf{T}^T, \dots, \gamma_n \mathbf{T}^T}_{3n}) \mathcal{M}, \end{aligned}$$

since $\text{vect}(\gamma \mathbf{P} \mathbf{T}) = \text{diag}(\gamma \mathbf{T}^T, \gamma \mathbf{T}^T, \gamma \mathbf{T}^T) \text{vect}(\mathbf{P})$. We deduce that the parameter vectors are related by $\check{\mathcal{X}} = \tilde{\mathbf{T}} \mathcal{X}$ where the *gauge transformation* $\tilde{\mathbf{T}}$ is defined by:

$$\tilde{\mathbf{T}} = \text{diag}(\underbrace{\gamma_1 \mathbf{T}^T, \dots, \gamma_n \mathbf{T}^T}_{3n}, \underbrace{\alpha_1 \mathbf{T}^{-1}, \dots, \alpha_m \mathbf{T}^{-1}}_m). \quad (5)$$

Note that $\det(\mathbf{T}) \neq 0 \Rightarrow \det(\tilde{\mathbf{T}}) \neq 0$ since the γ_i and α_j are non-zero. In general, \mathcal{X} and $\check{\mathcal{X}}$ are related by a unique gauge transformation.

The geometric equivalence of two reconstructions, denoted by \cong , is an equality ‘up to the gauge’ defined by:

$$(\mathcal{X} \cong \check{\mathcal{X}}) \Leftrightarrow (\exists \tilde{\mathbf{T}} \mid \check{\mathcal{X}} = \tilde{\mathbf{T}}\mathcal{X}). \quad (6)$$

with $\tilde{\mathbf{T}}$ of the form (5). We can now formally state the gauge invariance property. A bundle adjustment algorithm is gauge invariant if it preserves geometric equivalence: $(\mathcal{X} \cong \check{\mathcal{X}}) \Rightarrow ((\mathcal{X} + \delta) \cong (\check{\mathcal{X}} + \check{\delta}))$. By writing explicitly this property using equation (6), expanding, and since \mathcal{X} and $\check{\mathcal{X}}$ are related by a unique $\tilde{\mathbf{T}}$, we obtain the following definition of gauge invariance:

$$(\exists \tilde{\mathbf{T}} \mid (\check{\mathcal{X}} = \tilde{\mathbf{T}}\mathcal{X})) \Rightarrow (\check{\delta} = \tilde{\mathbf{T}}\delta). \quad (7)$$

Note that although derived in a different manner, this definition is similar to the one given in [12]. We examine successively under which conditions the Gauss-Newton and Levenberg-Marquardt iterations are gauge invariant.

4.2. Gauss-Newton

Consider a Gauss-Newton iteration. The increment δ within gauge \mathcal{G} is given by solving the normal equations (3). Due to gauge freedoms, matrix \mathbf{N} has a rank deficiency of g , and multiple solutions hold, corresponding to different gauges. Let \mathbf{G} be a $(p \times g)$ full column-rank matrix defined by $\mathbf{N}\mathbf{G} = \mathbf{0}$, i.e. the columns of \mathbf{G} span the nullspace of \mathbf{N} . Denoting by \dagger the Moore-Penrose pseudo-inverse, the possible solutions of the normal equations are parameterized by a $(g \times 1)$ vector \mathbf{v} as $\delta = -(\mathbf{J}^T\mathbf{J})^\dagger \mathbf{J}^T\mathbf{r} + \mathbf{G}\mathbf{v}$, where we substituted $\mathbf{N} = \mathbf{J}^T\mathbf{J}$ and $\mathbf{g} = \mathbf{J}^T\mathbf{r}$. Since $\mathbf{J} = \frac{\partial \mathbf{r}}{\partial \mathcal{X}} = \frac{\partial \mathbf{r}}{\partial \check{\mathcal{X}}} \frac{\partial \check{\mathcal{X}}}{\partial \mathcal{X}} = \tilde{\mathbf{J}}\tilde{\mathbf{T}}$, we obtain: $\delta = -(\tilde{\mathbf{T}}^T\tilde{\mathbf{J}}^T\tilde{\mathbf{J}}\tilde{\mathbf{T}})^\dagger \tilde{\mathbf{T}}^T\tilde{\mathbf{J}}\mathbf{r} + \tilde{\mathbf{T}}^{-1}\tilde{\mathbf{G}}\mathbf{v}$. From equation (7), gauge invariance holds if and only if:

$$(\tilde{\mathbf{T}}^T\tilde{\mathbf{N}}\tilde{\mathbf{T}})^\dagger = \tilde{\mathbf{T}}^{-1}\tilde{\mathbf{N}}^\dagger\tilde{\mathbf{T}}^{-T} \text{ and } \mathbf{v} = \check{\mathbf{v}},$$

which is verified if and only if $\tilde{\mathbf{T}}$ is orthonormal [12], i.e. $\tilde{\mathbf{T}}^{-1} = \tilde{\mathbf{T}}^T$. Hence, Gauss-Newton-based bundle adjustment is *not* gauge invariant. The previously derived condition does not leave enough flexibility to be exploited in this direction.

4.3. Levenberg-Marquardt

Consider a Levenberg-Marquardt iteration. The increment δ within gauge \mathcal{G} is given by solving the damped normal equations (4) $\mathbf{N}^*\delta = -\mathbf{g}$. By substituting $\mathbf{N}^* = \mathbf{N} + \mathbf{W}(\lambda) = \mathbf{J}^T\mathbf{J} + \mathbf{W}(\lambda)$ and $\mathbf{g} = \mathbf{J}^T\mathbf{r}$ in this equation and since $\det(\mathbf{N}^*) \neq 0$, we obtain: $\delta = -(\mathbf{J}^T\mathbf{J} + \mathbf{W}(\lambda))^{-1} \mathbf{J}^T\mathbf{r}$. By substituting $\mathbf{J} = \tilde{\mathbf{J}}\tilde{\mathbf{T}}$ and expanding, we get:

$$\begin{aligned} \delta &= -(\tilde{\mathbf{T}}^T\tilde{\mathbf{J}}^T\tilde{\mathbf{J}}\tilde{\mathbf{T}} + \mathbf{W}(\lambda))^{-1} \tilde{\mathbf{T}}^T\tilde{\mathbf{J}}\mathbf{r} \\ &= -\tilde{\mathbf{T}}^{-1}(\tilde{\mathbf{N}} + \tilde{\mathbf{T}}^{-T}\mathbf{W}(\lambda)\tilde{\mathbf{T}}^{-1})^{-1} \tilde{\mathbf{J}}\mathbf{r}. \end{aligned}$$

From equation (7), gauge invariance holds if and only if:

$$\check{\mathbf{W}}(\lambda) = \tilde{\mathbf{T}}^{-T}\mathbf{W}(\lambda)\tilde{\mathbf{T}}^{-1}. \quad (8)$$

The usual choices for matrices $\mathbf{W}(\lambda)$ and $\check{\mathbf{W}}(\lambda)$, e.g. $\mathbf{W}(\lambda) = \check{\mathbf{W}}(\lambda) = \lambda\mathbf{I}$, do not fulfill the above-derived property². Hence, standard implementations of Levenberg-Marquardt-based algorithms are *not* gauge invariant. Equation (8) can be verified if an appropriate gauge dependent choice is made for the damping matrices. This is the cornerstone for the gauge invariant method proposed in the next section.

5. Explicit and Implicit Gauges

Standard bundle adjustment algorithms are not gauge invariant. Hence, given a parameter estimate, there must exist an optimal local gauge within which the iteration reduces the reprojection error better than within the others. The problem of finding this gauge is dealt with in the next section. In this section, we examine the relationship between the damping matrix and gauge transformations.

A commonly used solution to globally enforce a gauge, e.g. [8, 11, 12], is what we call *explicit gauge* fixing. It consists in explicitly expressing the reconstruction within the desired gauge before each iteration. It implies the transfer of the points, the cameras, and all other entities being optimized, into the desired gauge.

We propose *implicit gauge* fixing. Consider equation (8). On the one hand, it tells us that the algorithm is not gauge invariant. On the other hand, it shows that there exists a close link between the gauge and the damping matrix. Hence, it can be used to perform bundle adjustment within gauge $\check{\mathcal{G}}$, while choosing $\check{\mathbf{W}}(\lambda)$ such that it behaves exactly as if it was conducted within gauge \mathcal{G} . For example if $\mathbf{W}(\lambda) = \lambda\mathbf{I}$ and $\check{\mathbf{W}}(\lambda) = \lambda\tilde{\mathbf{T}}^{-T}\tilde{\mathbf{T}}^{-1}$, then running the algorithm within gauges \mathcal{G} and $\check{\mathcal{G}}$ is strictly equivalent. We confirm in our experiments that explicit and implicit gauges give exactly the same results. Let $\mathbf{A} = \mathbf{T}^T\mathbf{T}$, then $\check{\mathbf{W}}(\lambda) = \lambda\tilde{\mathbf{T}}^{-T}\tilde{\mathbf{T}}^{-1}$ is given by:

$$\check{\mathbf{W}}(\lambda) = \lambda \text{diag}(\underbrace{\gamma_1^2\mathbf{A}^{-1}, \dots, \gamma_n^2\mathbf{A}^{-1}}_{3n}, \underbrace{\alpha_1^2\mathbf{A}, \dots, \alpha_m^2\mathbf{A}}_m). \quad (9)$$

In other words, a spherical trust region within gauge \mathcal{G} corresponds to a more complex trust region within gauge $\check{\mathcal{G}}$, encapsulated by the symmetric positive definite matrix $\check{\mathbf{W}}(\lambda)$ defined by equation (9).

To summarize, it is possible to reproduce the behaviour of bundle adjustment within a certain gauge, without explicitly expressing the reconstruction within this gauge, only by choosing the appropriate damping matrix, depending on the gauge transformation.

Based on this reasoning, we propose the following result.

²It can be shown that equation (8) with $\mathbf{W}(\lambda) = \check{\mathbf{W}}(\lambda) = \lambda\mathbf{I}$ is verified as in the Gauss-Newton case, i.e. if $\tilde{\mathbf{T}}$ is orthonormal.

Proposition 1 (Gauge invariance conditions) *Levenberg-Marquardt iterations can be made gauge invariant if the following two conditions are verified: (i) the form of the damping matrix is at least as general as the form given by equation (9) and (ii) the damping matrix is uniquely determined by a gauge dependent criterion defined such that equation (8) is verified.*

6. Gauge Dependent Damping

We propose a solution to choose a gauge dependent damping matrix $W(\lambda) = \lambda\mathcal{W}$ which has the three following important properties. First, it makes the underlying iteration gauge invariant. Second, it maximizes the decrease in the reprojection error up to first order. Third, it preserves the sparse block structure of the normal equations.

In order to achieve this result, we begin by deriving an approximation of the reprojection error, depending upon $W(\lambda)$. A quadratic approximation of $\mathcal{C}(\mathcal{X} + \delta)$ is given by equation (2). The incremental error e is given by:

$$e = \mathbf{g}^T \delta + \frac{1}{2} \delta^T \mathbf{N} \delta. \quad (10)$$

The increment vector δ is given by solving the augmented normal equations (4):

$$\delta = -(\mathbf{N} + W(\lambda))^{-1} \mathbf{g}.$$

Since $\det(W(\lambda)) \neq 0$, we can formulate the following Taylor expansion:

$$(\mathbf{N} + W(\lambda))^{-1} = \sum_{n=0}^{\infty} (-1)^n \mathbf{Z}(n),$$

where $\mathbf{Z}(n) = (W(\lambda)^{-1} \mathbf{N})^n W(\lambda)^{-1}$. This expression is valid provided $\|W(\lambda)^{-1} \mathbf{N}\| < 1$, i.e. when \mathbf{N} is small or λ is large. This leads to:

$$\delta = - \sum_{n=0}^{\infty} (-1)^n \mathbf{Z}(n) \mathbf{g}.$$

We substitute the above expression in equation (10). The first term is rewritten as:

$$\mathbf{g}^T \delta = - \sum_{n=0}^{\infty} (-1)^n \mathbf{g}^T \mathbf{Z}(n) \mathbf{g} = - \sum_{n=0}^{\infty} (-1)^n \mathbf{K}(n),$$

where $\mathbf{K}(n) = \mathbf{g}^T \mathbf{Z}(n) \mathbf{g}$. The second term becomes:

$$\frac{1}{2} \delta^T \mathbf{N} \delta = \frac{1}{2} \sum_{n=0}^{\infty} ((-1)^n \mathbf{g}^T \mathbf{Z}(n)) \mathbf{N} \sum_{n=0}^{\infty} ((-1)^n \mathbf{Z}(n) \mathbf{g}).$$

After expansion and using the following property arising from the symmetry of $W(\lambda)$: $\mathbf{Z}(n) \mathbf{N} \mathbf{Z}(n') = \mathbf{Z}(n+n'+1)$, we obtain:

$$\frac{1}{2} \delta^T \mathbf{N} \delta = \frac{1}{2} \sum_{n=0}^{\infty} (-1)^n (n+1) \mathbf{K}(n+1).$$

This gives the following Taylor expansion for e :

$$e = \sum_{n=0}^{\infty} (-1)^{n+1} \left(1 + \frac{n}{2}\right) \mathbf{K}(n). \quad (11)$$

Proposition 2 (Gauge invariant iterations) *The damping matrix defined such that equation (11) is minimized satisfies proposition 1. Hence, the underlying iteration is gauge invariant.*

Proof. We drop the parameter λ for this proof. Equation (11) can be rewritten as:

$$e = \mathbf{g}^T \left(-W^{-1} + \frac{3}{2} W^{-1} \mathbf{N} W^{-1} - \dots \right) \mathbf{g}.$$

Transfer the reconstruction from the gauge \mathcal{G} to $\tilde{\mathcal{G}}$ by applying the gauge transformation $\tilde{\mathbf{T}}$. By substituting $\mathbf{g} = \tilde{\mathbf{T}}^T \tilde{\mathbf{g}}$ and $\mathbf{N} = \tilde{\mathbf{T}}^T \tilde{\mathbf{N}} \tilde{\mathbf{T}}$, we obtain:

$$e = \tilde{\mathbf{g}}^T \tilde{\mathbf{T}} \left(-W^{-1} + \frac{3}{2} W^{-1} \tilde{\mathbf{T}}^T \tilde{\mathbf{N}} \tilde{\mathbf{T}} W^{-1} - \dots \right) \tilde{\mathbf{T}}^T \tilde{\mathbf{g}},$$

and hence, $\check{W}^{-1} = \tilde{\mathbf{T}} W^{-1} \tilde{\mathbf{T}}^T$, which verifies the gauge invariance equation (8) and concludes the proof. ■

Finding the damping matrix that minimizes e , equation (11), is a complicated problem. We propose to keep only the first-order term: $e \approx -\mathbf{g}^T W(\lambda)^{-1} \mathbf{g}$, which reduces the problem to finding a symmetric, positive definite, damping matrix $\lambda\mathcal{W}$ such that:

$$\mathcal{W} = \arg \max_{\mathcal{W}, \|\mathcal{W}\|^2=p} \mathbf{g}^T (\lambda\mathcal{W})^{-1} \mathbf{g}.$$

The normalization condition $\|\mathcal{W}\|^2 = p$ is the same as the standard choice $W(\lambda) = \lambda \mathbf{I}$ since $\|\mathbf{I}_{(p \times p)}\|^2 = p$. Obviously, the solution depends upon the form chosen for matrix \mathcal{W} . For example, if $\mathcal{W} = \text{diag}(\mathbf{d})$, where \mathbf{d} is a $(p \times 1)$ vector with strictly positive elements, then the problem transforms in:

$$\mathcal{W} = \text{diag} \left(\arg \max_{\mathbf{d}, \|\mathbf{d}\|^2=p} \sum_{k=1}^p g_k^2 / d_k \right),$$

which has the simple solution (up to scale) $\mathcal{W} \sim \text{diag}(\mathbf{g})^2$. According to proposition 1, this solution can not yield a subsequent gauge invariant algorithm since a diagonal damping matrix is not general enough for implicit gauges.

In a more general manner, since $W(\lambda)$ is a symmetric matrix, $W(\lambda)^{-1} = \frac{1}{\lambda} \mathbf{R}^T \mathbf{R}$, where \mathbf{R} is an upper-triangular matrix. With this parameterization, and since $\mathbf{g}^T \frac{1}{\lambda} \mathbf{R}^T \mathbf{R} \mathbf{g} = \frac{1}{\lambda} \|\mathbf{R} \mathbf{g}\|^2$, we obtain:

$$\mathbf{R} = \arg \max_{\mathbf{R}, \|\mathbf{R}^T \mathbf{R}\|^2=p} \frac{1}{\lambda} \|\mathbf{R} \mathbf{g}\|^2. \quad (12)$$

Hence, the coefficients of matrix \mathbf{R} can be found by solving a simple linear least squares optimization problem.

1. Compute the damping matrix $\mathcal{W}(\lambda)^{-1} = \lambda \mathbf{R}^T \mathbf{R}$ where:

$$\mathbf{R} = \text{diag}(\underbrace{\mathbf{M}_1, \dots, \mathbf{M}_n}_{3n}, \underbrace{\mathbf{S}, \dots, \mathbf{S}}_m).$$

The (4×4) upper-triangular matrices \mathbf{M}_i and \mathbf{S} are formed by solving (see main text):

$$\mathbf{M}_i = \arg \max_{\mathbf{M}_i, \|\mathbf{M}_i^T \mathbf{M}_i\|^2=12} \sum_{k=1}^3 \|\mathbf{M}_i \mathbf{g}_{\mathcal{M},i}^k\|^2$$

$$\mathbf{S} = \arg \max_{\mathbf{S}, \|\mathbf{S}^T \mathbf{S}\|^2=4} \sum_{j=1}^m \|\mathbf{S} \mathbf{g}_{\mathcal{S},j}\|^2,$$

where $\mathbf{g}_{\mathcal{S},j}$ and $\mathbf{g}_{\mathcal{M},i}^k$ are (4×1) gradient vectors for the j -th point and for the k -th row of the i -th camera matrix respectively.

2. Perform one Levenberg-Marquardt iteration with the damping matrix $\mathcal{W}(\lambda)$. This iteration is gauge invariant due to the above choice.
3. Optional: Project the estimate onto a global gauge.

Table 1. The gauge invariant iteration we propose. Periodical enforcement of global gauge constraints (step 3) such as renormalization of homogeneous coordinates is recommended. Gauge invariance holds in the sense that any global gauge can be enforced in step 3 without affecting the result.

From proposition 1 and imposing the fact that the damping should not spoil the sparse block structure of the normal equations, we are left with the following possible form:

$$\mathbf{R} = \text{diag}(\underbrace{\mathbf{R}_{\mathcal{M},1}, \dots, \mathbf{R}_{\mathcal{M},n}}_n, \underbrace{\mathbf{R}_{\mathcal{S},1}, \dots, \mathbf{R}_{\mathcal{S},m}}_m),$$

where the $\mathbf{R}_{\mathcal{M},i}$ and the $\mathbf{R}_{\mathcal{S},j}$ are respectively (12×12) and (4×4) upper-triangular matrices. Given that each of the p parameters gives one constraint on the damping matrix and since it must be uniquely determined (i.e. it must not have more than p parameters), we propose the following $10(n+1)$ -parameter choice:

$$\mathbf{R}_{\mathcal{M},i} = \text{diag}(\mathbf{M}_i, \mathbf{M}_i, \mathbf{M}_i) \text{ and } \mathbf{R}_{\mathcal{S},j} = \mathbf{S},$$

where the \mathbf{M}_i and \mathbf{S} are (4×4) upper-triangular matrices. This choice means that the parameter block of each camera i has its own 10 parameter damping block $\mathbf{M}_i^T \mathbf{M}_i$, while a unique 10 parameter damping block $\mathbf{S}^T \mathbf{S}$ is defined for point parameters. This form is handy since each block can be found independently by solving a problem of the form (12). More details are given in table 1. Below, we

give details about solving the low-dimensional maximization problems of table 1. Let the problem to be solved be $\max_{\mathbf{x}, \|\mathbf{x}\|^2=b} \|\mathbf{C}\mathbf{x}\|^2$. The singular vector \mathbf{v} associated to the largest singular value of matrix \mathbf{C} gives the solution for \mathbf{x} as $\mathbf{x} = \sqrt{b}\mathbf{v}$. Singular value decomposition of matrix \mathbf{C} can be used to compute \mathbf{v} .

7. Experimental Results

We compare our gauge invariant algorithm, denoted GAUGE_INVARIANT, to some other ones. FREE directly optimizes the camera matrices and 3D points. Pseudo-inverse is used to solve the normal equations. HARTLEY [6] consists in using a reference camera to partially eliminate the gauge and BARTOLI [2] consists in using two reference cameras to eliminate the gauge. FAUGERAS [4] is based on fixing the coordinates of 5 reference points to enforce the gauge. A similar approach is proposed by Mohr et al in [13]. MCLAUCHLAN [11] consists in enforcing a normalized basis after each iteration and in using first order gauge constraints in the normal equations and KANATANI [8] is similar to method FREE but periodically projects the estimate on a pre-chosen gauge to prevent it to drift too much. When reference cameras or points are required by the parameterization, they are chosen at random. Note that we have implemented all these algorithms in a unified manner: the optimization engine is unique, only the parameterization and the damping strategy change.

We measure the number of iterations and the reprojection error at convergence. The initial solution, denoted by INIT, is computed by registering cameras in turn to a two-view reconstruction. Image point coordinates are standardized such that they lie in $[-1 \dots 1]$.

7.1. Simulated Data

We simulate $m = 100$ points lying in a cube with 1 meter side length, observed by $n = 5$ cameras with a focal length of 1000 pixels. The points are offset from a base plane lying inside the cube, with a mean offset denoted by d . Cameras are situated 10 meters away from the center of the cube. The baseline between consecutive cameras is 3 meters. All points are visible in all views. We add a centered Gaussian noise on true point positions with a 1 pixel variance. We vary some parameters of the above-described setup to compare the algorithms in different situations. The results are averaged over 50 trials.

Figure 1 shows the results when the scene flatness, i.e. the mean offset d from the plane, is varied. We observe that the reprojection errors, i.e. the accuracy, are undistinguishable for all methods, except for methods FREE which converges to a different local minimum than the others, for weak geometry, i.e. when the mean offset from the plane d is small. Concerning the number of iterations, i.e. the computational cost, we observe that when the mean offset is large, i.e. when the geometry is strong, there is only slight differences between the different methods besides for method FREE, which

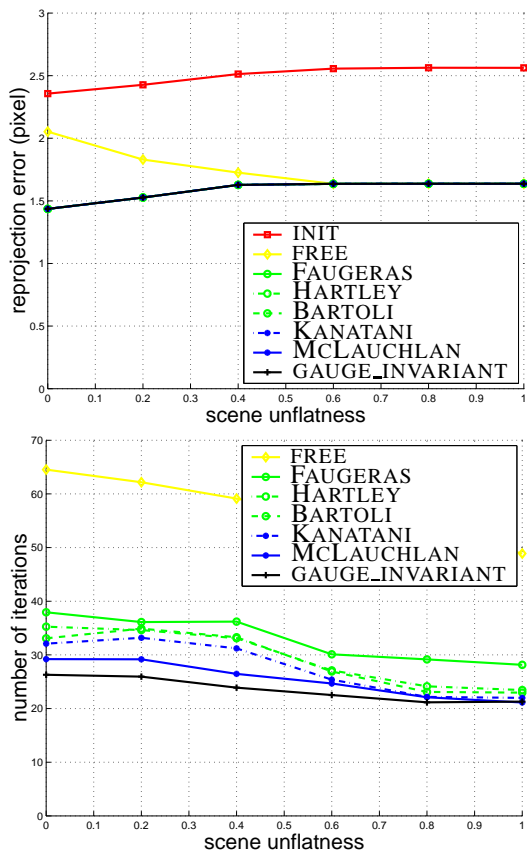


Figure 1. Reprojection error and number of iterations when varying the scene unflatness, from weak to strong geometry.

takes clearly more iterations to converge. As the geometry becomes weaker, i.e. when the offset decreases, large discrepancies can be observed between the different methods. Method FAUGERAS based on reference points gives bad results, since using reference points to fix a projective basis can be very unstable. Methods HARTLEY and BARTOLI based on reference cameras and KANATANI based on periodical projection on a pre-chosen gauge give reliable results when the geometry is strong enough. Method MCLAUCHLAN performs better, while GAUGE_INVARIANT is the method the less sensitive to the instability of the scene.

We tested other strong to weak scene configurations, based on varying the number of points and cameras, the baseline between consecutive cameras and the visibility (not shown here due to lack of space). We observed similar results as above: strong geometries reduce the discrepancies between the different methods. Similarly, we tested the influence of the initialization. As expected, discrepancies between the different methods reduce as the initialization gets more accurate.

Figure 2 shows the reprojection error as a function of the iterations for the scene setting with $d = 0.5$. As expected, we observe that the error is faster to decrease as the number of iterations gets low.

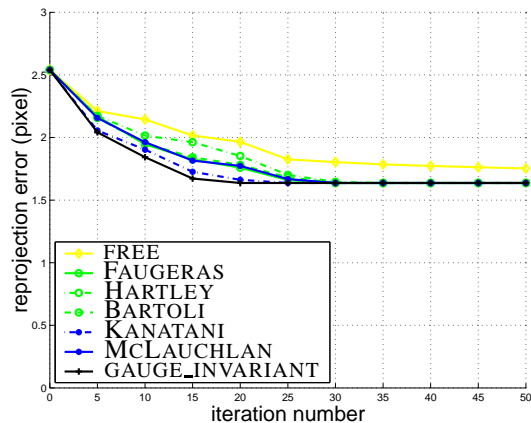


Figure 2. Reprojection error as a function of the iterations.

7.2. Real Data

We compare the algorithms on various image streams. For two of them, the office sequence, figure 3, and the hotel sequence³, we show results, see table 2. For the office sequence, all methods give a reprojection error of 0.8392 pixel. For the hotel sequence, they all gave 0.9594 pixel, besides method FREE which gave 1.5268 pixels. Concerning the number of iterations, the same observations as for simulated data can be made. The hotel sequence gives a weak geometry since an affine camera model is well-adapted to these data. Hence, a full perspective projection model is not well-constrained, which tends to increase the discrepancies between the different methods.

Algorithm	office seq.	hotel seq.
FREE	12	28
HARTLEY	7	16
BARTOLI	8	19
FAUGERAS	10	22
MCLAUCHLAN	6	15
KANATANI	7	16
GAUGE_INVARIANT	7	13

Table 2. Number of iterations of the algorithms for the office and the hotel sequences.

³These data have been provided by the Modeling by Videotaping group in the Robotics Institute, Carnegie Mellon University.




Figure 3. One out of the 15 frames of the office sequence, overlaid with the 38 corner features (left) and snapshot of the reconstructed shape (right).

8. Conclusions

We proposed a bundle adjustment algorithm which maximizes the decrease in the error at each iteration, regardless the gauge within which the reconstruction is expressed.

We derived this algorithm based on a careful study of how the gauge influence the iterations of standard algorithms, which are not gauge invariant. In particular, concerning Levenberg-Marquardt iterations, we showed that the standard diagonal damping matrices, defining elliptical trust regions, transform to a more complex shaped trust region when changing the gauge. Based on this, we proposed a gauge dependent damping matrix and a practical algorithm to compute it, that allows gauge invariant iterations while maximizing the decrease in the reprojection error. The final algorithm can be incorporated in a straightforward manner to existing minimization engines since it consists in appropriately choosing the damping matrix by solving low-dimensional linear least squares systems. The sparse block structure of the normal equations is preserved.

We compared our algorithm to existing ones using simulated and real data. We observed that it is more reliable in the sense that it is less sensitive to weak geometry and weak initialization. Moreover, the error decreases more intensely throughout the iterations. Most algorithms converge to the same solution, but require different number of iterations.

References

- [1] K. Atkinson, editor. *Close Range Photogrammetry and Machine Vision*. Whittles Publishing, 1996.
- [2] A. Bartoli. On the non-linear optimization of projective motion using minimal parameters. In *Proceedings of the 7th European Conference on Computer Vision, Copenhagen, Denmark*, volume 2, pages 340–354, May 2002.
- [3] A. Dermanis. The photogrammetric inner constraints. *ISPRS Journal of Photogrammetry and Remote Sensing*, 49(1):25–39, 1994.
- [4] O. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In G. Sandini, editor, *Proceedings of the 2nd European Conference on Computer Vision*,

Santa Margherita Ligure, Italy, pages 563–578. Springer-Verlag, May 1992.

- [5] R. Hartley. Euclidean reconstruction from uncalibrated views. In *Proceeding of the DARPA-ESPRIT workshop on Applications of Invariants in Computer Vision, Azores, Portugal*, pages 187–202, October 1993.
- [6] R. Hartley. Projective reconstruction and invariants from multiple images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):1036–1041, October 1994.
- [7] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, June 2000.
- [8] K. Kanatani and D. D. Morris. Gauges and gauge transformations for uncertainty description of geometric structure with indeterminacy. *IEEE Transactions on Information Theory*, 47(5), July 2001.
- [9] K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, pages 164–168, 1944.
- [10] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, June 1963.
- [11] P. F. McLauchlan. Gauge invariance in projective 3D reconstruction. In *Proceedings of the Multi-View Workshop, Fort Collins, Colorado, USA*, 1999.
- [12] P. F. McLauchlan. Gauge independence in optimization algorithms for 3D vision. In *Proceedings of the Vision Algorithms Workshop, Dublin, Ireland*, 2000.
- [13] R. Mohr, L. Quan, and F. Veillon. Relative 3D reconstruction using multiple uncalibrated images. *The International Journal of Robotics Research*, 14(6):619–632, 1995.
- [14] D. D. Morris, K. Kanatani, and T. Kanade. Gauge fixing for accurate 3d estimation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition, Kauai, Hawaii, USA*, December 2001.
- [15] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C - The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.
- [16] G. A. F. Seber and C. J. Wild. *Non-linear regression*. John Wiley & Sons, New York, 1989.
- [17] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle adjustment — a modern synthesis. In B. Triggs, A. Zisserman, and R. Szeliski, editors, *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice, Corfu, Greece*, volume 1883 of *Lecture Notes in Computer Science*, pages 298–372. Springer-Verlag, 2000.