



HAL
open science

Edge-based multi-modal registration and application for night vision devices

Camille Sutour, Jean-François Aujol, Charles-Alban Deledalle, Baudouin Denis de Senneville

► **To cite this version:**

Camille Sutour, Jean-François Aujol, Charles-Alban Deledalle, Baudouin Denis de Senneville. Edge-based multi-modal registration and application for night vision devices. *Journal of Mathematical Imaging and Vision*, 2015, pp.1-20. 10.1007/s10851-014-0554-8 . hal-01068431v2

HAL Id: hal-01068431

<https://hal.science/hal-01068431v2>

Submitted on 2 Apr 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Edge-based multi-modal registration and application for night vision devices

Camille Sutour · Jean-François Aujol · Charles-Alban Deledalle · Baudouin Denis de Senneville

Abstract Multi-modal image sequence registration is a challenging problem that consists in aligning two image sequences of the same scene acquired with a different sensor, hence containing different characteristics. We focus in this paper on the registration of optical and infra-red image sequences acquired during the flight of a helicopter. Both cameras are located at different positions and they provide complementary informations. We propose a fast registration method based on the edge information: a new criterion is defined in order to take into account both the magnitude and the orientation of the edges of the images to register. We derive a robust technique based on a gradient ascent and combined with a reliability test in order to quickly determine the optimal transformation that matches the two image sequences. We show on real multi-modal data that our method out-

performs classical registration methods, thanks to the shape information provided by the contours. Besides, results on synthetic images and real experimental conditions show that the proposed algorithm manages to find the optimal transformation in few iterations, achieving a rate of about 8 frames per second.

Keywords Multi-Modal · image sequence registration · night vision · optimization

1 Introduction

1.1 Operational context

Multi-modal image registration consists in aligning several images of a same scene acquired by different sensors, from a different point of view or at a different time. It is widely used in medical applications, for example for comparing images of the brain obtained with computer tomography (CT) to positron emission tomography (PET) or magnetic resonance (MRI) images. Multi-modal registration is also studied in remote sensing applications, for example, for the association of a synthetic aperture radar (SAR) image and an optical one. It is an important preliminary step for high level analysis such as image fusion, change detection, augmented reality, etc. A survey of most registration methods can be found in [4, 27].

The goal of this paper is to perform multi-modal registration between optical image sequences obtained from a night vision device and infra-red image sequences, acquired from a helicopter, in the perspective of fusing both modalities. Optical images are obtained thanks to a

C. Sutour thanks the DGA and the Aquitaine region for funding her PhD. J.-F. Aujol acknowledges the support of the Institut Universitaire de France. This study has been carried out with financial support from the French State, managed by the French National Research Agency (ANR) in the frame of the "Investments for the future" Programme IdEx Bordeaux - CPU (ANR-10-IDEX-03-02).

C. Sutour
IMB and LaBRI, Université de Bordeaux, Talence, France,
E-mail: camille.sutour@math.u-bordeaux.fr

J.-F. Aujol, Ch.-A. Deledalle, B. Denis de Senneville
IMB, CNRS, UMR 5251, Université de Bordeaux, Talence, France,
E-mail: {jaujol,cdeledal,bdenisde}@math.u-bordeaux.fr

B. Denis de Senneville
Imaging Division, UMC Utrecht

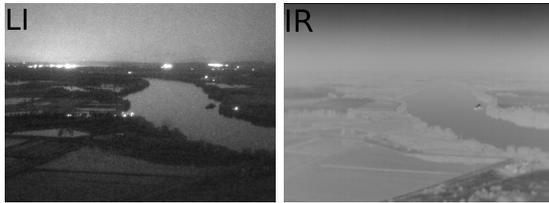


Fig. 1 Example of (left) optical (LI) and (right) infra-red (IR) images simultaneously acquired from a flying helicopter. The resolution and the information displayed are different, as well as the intensity distributions.

light intensifier (LI) that multiplies the photons in order to amplify the luminosity. The light intensifier is combined with a CCD camera to obtain numerical images. The LI device is located on the helmet of the pilot and the images are projected on the visor. They display natural scenes and are easily interpretable, but they suffer from classical defaults inherent to night vision devices: they are degraded by (photon count) noise and they suffer from artifacts (meshing, changes of illumination...), they are poorly contrasted while using a large dynamic range, and they are saturated around light sources. An example is shown in Figure 1. On the other hand, infra-red (IR) images reflect the temperature of the scene. They are not easy to interpret because they do not reflect the intuitive perception of the scene, as shown on Figure 1. However, they provide precious information such as vehicles, roads and buildings because they are hot sources compared to the ground. The infra-red camera is located at the bottom of the helicopter, and can be driven by the pilot. These two video cameras observe the scene from a different angle, and they can move independently from each other, so a careful registration that takes into account both the difference of perspective and the relative movement between the two must be achieved prior to combining the information.

In the scope of this study, the optical (LI) image and the infra-red (IR) image are acquired simultaneously with the same update rate. Each time a new couple of images (LI, IR) is obtained, the goal is to register the optical image into the frame of reference of the infra-red image IR. This consists in finding the global spatial transformation T that associates each pixel of the LI image (referred to as the current image u) to its corresponding location in the IR image (referred to as the reference image v).

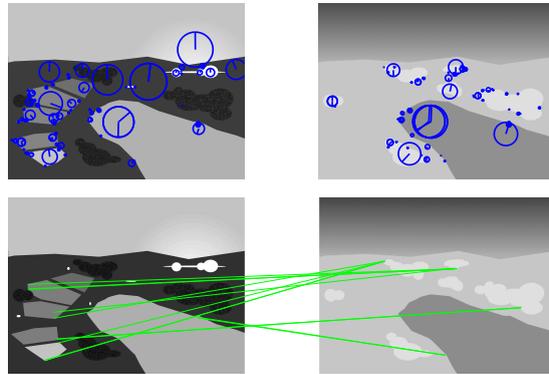


Fig. 2 Detection of SIFT features and associated descriptors on each modality, and matching of the features obtained from different modalities.

1.2 State of the art

Registration techniques mostly stem from medical applications, remote sensing, or computer vision. In computer vision, registration techniques are often based on the detection and matching of special features. Among them, the scale invariant feature transform (SIFT) descriptor proposed by Lowe in [13] combines a scale invariant region detector with a descriptor based on the gradient distribution in the detected regions. These descriptors have been widely studied and adapted [3, 16], combined with classical matching algorithms such as the RANSAC algorithm [7], in order to estimate the optimal transformation. The SIFT descriptors have also been adapted for the registration of CT and MR volumes in [1, 24, 18] or for SAR images in [6], but the proposed SIFT implementations cannot intrinsically deal with a multi-modal framework. Indeed, even if each modality benefits from its own feature descriptors, it is not always possible to match the descriptors together due to the different information inherent to each modality [26]. Figure 2 shows on synthetic images the difficulty to associate features extracted from different sensors: even though each modality has its own features, the descriptors do not allow to perform accurate feature matching, so no transformation can be directly estimated.

An alternative to deal with different modalities consists in finding the optimal domain transformation T that maximizes a given energy. The cross-correlation metric [20, 21] measures the correlation between the values of the current transformed image, denoted $u(T)$, and

the reference one v using the following formula:

$$CC(T) = \frac{\int_{\Omega} u_0(T(X)) \cdot v_0(X) dX}{\sqrt{\int_{\Omega} u_0(T(X))^2 dX} \cdot \sqrt{\int_{\Omega} v_0(X)^2 dX}}. \quad (1)$$

where Ω is the (continuous) image domain, and $u_0(T(X))$ and $v_0(X)$ are the centered values of the images, i.e. the difference between the image and the average value of the image $\overline{u(T)}$ or \overline{v} : $u_0(T(X)) = u(T(X)) - \overline{u(T)}$ and $v_0(X) = v(X) - \overline{v}$. The squared cross-correlation [11] can also be more suitable for multi-modal registration, but both metrics are based on the assumption that the intensities of the images to register are close up to an affine scaling, which is not the case in our multi-modal problem.

Mutual information [25, 14] reflects the relation between the intensity distributions, but without any assumption regarding the nature of this relation. Mutual information is issued from information theory and the notion of entropy. It can also be interpreted in terms of a Kullback-Leibler distance:

$$I(T) = \int_{\Omega^2} p(u(T(X)), v(Y)) \cdot \log \frac{p(u(T(X)), v(Y))}{p(u(T(X)))p(v(Y))} dXdY, \quad (2)$$

where $p(u, v)$ is the joint probability distributions of u and v and $p(u)$ and $p(v)$ are the marginal distributions. It measures the information that one data contains about the other: the more independent u and v are, the closer the joint probability $p(u, v)$ is to the distribution $p(u) \times p(v)$. In practice, these distributions are estimated by computing the marginal and joint histograms of the values of the images to register. The registration is then performed by seeking the transformation that will maximize the mutual information between both images. However, it still requires that the intensities of both modalities are close to be in bijection. Unfortunately, this is not satisfied in our problem: some areas are highly textured on the optical image but smooth on the IR image, while some constant areas on the optical image (such as the sky or the river on Figure 1) are shaded on the IR image.

Registration can also be performed using an appropriate image representation. In [12], the authors define an image representation based on four directional derivative filters that extract the edge information along four directions, then

the registration is performed simultaneously on the four directions using a local-normalized-correlation. The use of the edge information as a relevant image representation has then been widely adopted. In [23], an edge-based metric is defined in order to measure the alignment of the gradient ∇v of the reference image and the gradient $\nabla u(T)$ of the transformed version of the image to register, where T is the tested transformation. The edge alignment is evaluated at each pixel thanks to the following edge-based criterion:

$$C_S(T) = \frac{\int_{\Omega} w_T(X) \cos(2\Delta\theta_T(X)) dX}{\int_{\Omega} w_T(X) dX} \quad (3)$$

where $w_T(X)$, $\Delta\theta_T(X)$ are based on the magnitude M and the orientation θ of the image gradients at location X :

$$\begin{aligned} w_T(X) &= M_u(T(X))M_v(X), \\ \Delta\theta_T(X) &= \theta_u(T(X)) - \theta_v(X). \end{aligned} \quad (4)$$

The $\cos(2\Delta\theta_T(X))$ in equation (3) favors the transformations that align the edges, regardless of the contrast reversals, i.e. gradient direction. Besides, when dealing with multi-modal images, some discontinuities can only appear in one of the two modalities, so the weight $w_T(X)$ favor strong edges that occur in both modalities.

In [8], a similar edge-based metric is used, based on the following quantity:

$$\omega_T(X) = \left\langle \frac{\nabla u(T(X))}{\|\nabla u(T(X))\|_{\epsilon}}, \frac{\nabla v(X)}{\|\nabla v(X)\|_{\epsilon}} \right\rangle^2 \quad (5)$$

where $\|\nabla v(X)\|_{\epsilon} = \sqrt{\nabla v(X)^T \nabla v(X) + \epsilon^2}$, and ϵ is automatically selected based on the noise variance estimation in order to distinguish between edges and noise fluctuations, even though it can be hard to perform on operational data. The Normalized Gradient Fields (NGF) metric is then defined as follows :

$$\begin{aligned} C_H(T) &= \int_{\Omega} \omega_T(X) dX \\ &= \int_{\Omega} \cos^2(\Delta\theta_T(X)) dX \end{aligned} \quad (6)$$

This metric uses normalized gradients; it can be expressed as a scalar product or as the cosine of the angle between the edges, regardless of the edge amplitude.

This edge-based metric can be traced back to [19] where the shape information is combined with the mutual information. In order to take into account the edges that appear in

both modalities, the scalar product $\omega_T(X)$ is weighted by the minimum of the gradient magnitude:

$$G(T) = \int_{\Omega} \omega_T(X) \min(|\nabla u(T(X))|, |\nabla v(X)|) \, dX, \quad (7)$$

then $G(T)$ is combined with mutual information in order to take into account both spatial and distribution-based information:

$$\text{MI-G}(T) = I(T) \times G(T). \quad (8)$$

Finally, in [10], the authors have developed modality invariant local image descriptors based on the notion of self-similarity derived from the NL-means algorithm [5]. They associate to each pixel X of the image u a descriptor based on the local similarities:

$$\text{MIND}(u, X, r) = \frac{1}{n} \exp\left(-\frac{D_p(u, X, X+r)}{V(u, X)}\right), \quad (9)$$

where n is a normalization constant, $r \in R$ defines a search region in which the patches are compared, D_p is the distance between patches of size p , usually a pixel-wise square distance, and $V(u, X)$ is an estimation of the local variance in order to take into account the noise. This MIND descriptor associates a vector of size $|R|$ to each pixel of the image u , then the descriptors of each modality are compared using the following metric:

$$S(X, T) = \frac{1}{|R|} \sum_{r \in R} |M_u - M_v|, \quad (10)$$

with $M_{u(T)} = \text{MIND}(u(T), X, r)$ and $M_v = \text{MIND}(v, X, r)$. Minimizing the summed distance $\sum_X S(X, T)$ returns the best transformation T that puts in correspondence the similarity information between the patches of both modalities, hence being sensitive to the structural information rather than to the intensities. However, this metric is prone to the same limitations than the non-local methods: it relies on several parameters such as the local variance estimation or the size of the patches and the search region, and it is affected by low contrasts. Besides, if the structural information is different from one modality to the other, for example on an area that is smooth on one image and textured on the other, then the descriptors cannot be compared reliably.

1.3 Contribution and organization of the paper

We propose to extend the edge-based metrics of (3) and (6) to a robust night vision framework as follows: we develop a new criterion that takes into account both the magnitude and the direction of the edges, and we maximize this criterion using a gradient ascent scheme in order to find the best transformation that will align one image with the other.

Our main contributions are the new criterion we propose, and the gradient ascent optimization combined with a temporal validation scheme. The theoretical and experimental study we conduct validate the proposed model, and we show that the algorithm can proceed up to 8 images per second, which makes it suitable to an embedded operational registration.

The proposed model is presented in section 2, then section 3 provides an optimization scheme based on a gradient ascent, and combined with a temporal scheme that guarantees error control and robustness. Section 4 studies the performance of the metric: we show that the maximization of the proposed criterion does allow to recover the optimal transformation, both in theory and in practice, and we study the stability and robustness of this metric. We also check that the gradient ascent scheme allows to recover the optimal transformation parameters. Section 5 presents an extension of the proposed model to the general case of projective transformations and extends the gradient ascent accordingly. Finally, section 6 shows results on real data.

2 Multi-modal framework

2.1 Definition of the criterion

The current image u is registered to the reference position given by v as follows. Let M_v and θ_v be the magnitude and the orientation of the gradient of the reference image v , computed using a Sobel edge detector [22]. M_u and θ_u are defined similarly in the image to register u .

We define an edge-based metric that is adapted to the night vision framework and the characteristics of each modality, and that is easy to manipulate in an embedded operational context. We define the following criterion:

$$C(T) = \int_{\Omega} |\nabla u(T(X)) \cdot \nabla v(X)| \, dX, \quad (11)$$

that can also be written under the following form:

$$C(T) = \int_{\Omega} w_T(X) |\cos(\Delta\theta_T(X))| dX, \quad (12)$$

where $w_T(X)$ and $\Delta\theta_T(X)$ refer to the magnitude and orientation quantities defined in (4). This criterion favors strong edges, thanks to the amplitude weighting, and it is insensitive to contrast reversals, i.e. opposite directions of the gradient: thanks to the absolute value of the cosine, both parallel and anti-parallel edges are considered to coincide. It allows to take into account edges that occur in both modalities, regardless of the direction of the intensity changes. Besides, contrary to the metric proposed in [23], this criterion is not normalized. This makes it easier to manipulate (more stable), and it is more sensitive to the number of edges that are actually put in correspondence. Indeed, the normalized criterion performs a weighted average of the score obtained for each edge, so it measures the average edge alignment that has been performed on all the edges that occur in both modalities. On the contrary, this un-normalized criterion adds up the score of each aligned edge, so that the more edges are in correspondence the higher the criterion is. The normalized criterion might favor very precise alignments, regardless of the number of matches, while being sensitive to mismatches, whereas this un-normalized criterion might prefer slightly less precise matches, if they occur often enough.

2.2 Transformation model

In the original paper of Sun et al. [23], the criterion of (3) is optimized by performing an exhaustive search on all the possible transformation parameters, that originally consist of a translation in both directions. In the scope of our application, we have first considered for possible transformations a translation in both directions (horizontal and vertical), and a uniform zoom. We denote the zoom parameter z , and the translation parameters in the horizontal and vertical direction respectively t_1 and t_2 . If $X = (x \ y \ 1)^T$ are the coordinates of the image to register (that we can also note in the concise form $X = (x \ y)^T$), we can define the

transformation matrix $T = T_{t_1, t_2, z}$ as :

$$\begin{aligned} T_{t_1, t_2, z}(X) &= \begin{pmatrix} 1 & 0 & t_1 \\ 0 & 1 & t_2 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1+z & 0 & 0 \\ 0 & 1+z & 0 \\ 0 & 0 & 1 \end{pmatrix} X \\ &= \begin{pmatrix} 1+z & 0 & t_1 \\ 0 & 1+z & t_2 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \end{aligned} \quad (13)$$

3 Proposed optimization scheme

3.1 Gradient ascent

Thanks to the formulation proposed in equation (11), an explicit optimization scheme is derived to maximize the proposed metric at each iteration n , by performing a gradient ascent on the transformation $T_{t_1, t_2, z}$:

$$\begin{cases} t_1^{n+1} = t_1^n + \lambda_1 \partial_{t_1} C(T_{t_1^n, t_2^n, z^n}) \\ t_2^{n+1} = t_2^n + \lambda_2 \partial_{t_2} C(T_{t_1^n, t_2^n, z^n}) \\ z^{n+1} = z^n + \lambda_3 \partial_z C(T_{t_1^n, t_2^n, z^n}) \end{cases}, \quad (14)$$

where the derivatives of the function $C(T_{t_1, t_2, z})$ are at each iteration:

$$\begin{aligned} \partial_{t_1} C(T_{t_1, t_2, z}) &= \int_{\Omega} \sigma D^2 u(T_{t_1, t_2, z}(X)) \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \cdot \nabla v(X) dX, \\ \partial_{t_2} C(T_{t_1, t_2, z}) &= \int_{\Omega} \sigma D^2 u(T_{t_1, t_2, z}(X)) \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \cdot \nabla v(X) dX, \\ \partial_z C(T_{t_1, t_2, z}) &= \int_{\Omega} \sigma D^2 u(T_{t_1, t_2, z}(X)) \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \cdot \nabla v(X) dX. \end{aligned} \quad (15)$$

where $\sigma = \text{sign}(\nabla u(T_{t_1, t_2, z}(X)) \cdot \nabla v(X))$.

The computation of the derivatives is detailed in appendix A. The functional we seek to maximize is subject to local maxima, so the initialization is important. For the first frame, we can either perform a coarse exhaustive search as in [23], or perform several gradient ascents with different initializations and select the result that gives the best metric value. Then in practice, the sequence provides temporal regularity, so the transformation for each frame can be initialized with the parameters obtained from the previously acquired frame.

3.2 Temporal implementation

In order to accelerate the convergence, improve the liability of the algorithm and control the performance of the registration, the gradient ascent has been included into a temporal scheme that uses the information from the previously

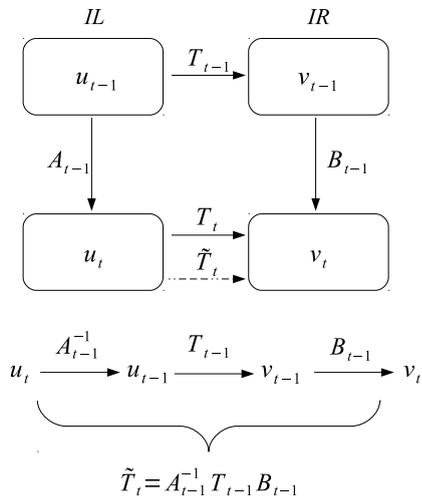


Fig. 3 Proposed temporal scheme for a fast convergence of the gradient ascent algorithm, and error control.

registered frames to predict and control the registration for the next frames. Figure 3 displays the different steps for a fast, in-flight registration. At time $t-1$, the LI and IR images respectively called u_{t-1} and v_{t-1} are registered with transformation T_{t-1} , so that:

$$u_{t-1}(T_{t-1}X) = v_{t-1}(X). \quad (16)$$

Besides, for a single modality between u_{t-1} and u_t or v_{t-1} and v_t , movement can be directly estimated in-flight. Indeed, the helmet is equipped with a posture detection system that provides information about the movements of the head, while the inertial systems can assess the general movements of the helicopter. This allows a fast motion estimation at no greater cost.

We perform this temporal registration separately on each modality, in order to estimate the transformations A_{t-1} between u_{t-1} and u_t and B_{t-1} between v_{t-1} and v_t , such that:

$$\begin{cases} u_{t-1}(A_{t-1}X) = u_t(X) \\ v_{t-1}(B_{t-1}X) = v_t(X) \end{cases} \quad (17)$$

Thanks to these three estimations, it is possible to predict the estimated transformation \tilde{T}_t between u_t and v_t as:

$$\tilde{T}_t = A_{t-1}^{-1} T_{t-1} B_{t-1} \quad (18)$$

This estimation \tilde{T}_t can be used as a close enough initialization at time t . This ensures that the gradient ascent will converge in few iterations, and since the initialization is reasonable it will lead to a relevant maximum.

Besides, this procedure can also be used to control the energy and prevent any divergence of the gradient ascent process. Indeed, the algorithm can be subject to local maxima and it is quite sensitive to the gradient steps for each parameter: if the optimum is close and the gradient steps are too large, any update might overshoot and go over the maximum. The temporal smoothness assumption can balance this sensitivity: if the energy at the end of the gradient ascent is found to be lower than the initialization, this means that the algorithm has not converged properly or that it has been sent to a local maximum, so we can choose to stick to the estimation \tilde{T}_t .

4 Analysis and validation of the proposed model

4.1 Theoretical analysis

The goal of this section is to study the registration from a mathematical point of view in order to show that the maximization of the proposed criterion does result theoretically in finding the optimal parameters.

We study the one dimensional case and we focus on aligning two edges when the signal to register is subject to a translation and a zoom. The registration of only one edge is subject to an aperture problem, since an edge, seen from different levels of zoom, remains the same. To remedy this, we use for a reference signal a box function $v_0(x) = 1$ if $x \in [-1, 1]$, 0 otherwise (see Figure 4).

The signal to register is then defined as $u_0(x) = v_0(ax - b)$ where a is the zoom parameter with $a > 0$ (a corresponds to the factor $1 + z$ in the transformation model described in section 2, equation (13)) and b the translation parameter.

Intuition

The reference signal described above is not differentiable in ± 1 . However, its derivative ∇v_0 can be represented as the sum of two Diracs at location ± 1 : $\nabla v_0(x) = \delta_{-1} - \delta_1$. We can also define $\nabla u_0(x) = a \cdot \nabla v(ax - b) = a \left(\delta_{\frac{-1+b}{a}} - \delta_{\frac{1+b}{a}} \right)$, so that the functional that we seek to maximize can be expressed in a

heuristic way as:

$$\begin{aligned}
F(a, b) &= \int_{\mathbb{R}} |\nabla u_0(x) \cdot \nabla v_0(x)| dX \\
&= \int_{\mathbb{R}} a |\nabla v_0(ax - b) \cdot \nabla v_0(x)| dX \\
&= \int_{\mathbb{R}} a \left| \left(\delta_{\frac{-1+b}{a}} - \delta_{\frac{1+b}{a}} \right) \cdot (\delta_{-1} - \delta_1) \right| dX \\
&= \int_{\mathbb{R}} a \left(\delta_{\frac{-1+b}{a}} + \delta_{\frac{1+b}{a}} \right) \cdot (\delta_{-1} + \delta_1) dX
\end{aligned} \tag{19}$$

Although it is not formally correct to deal with Dirac products, one can presume how the functional is going to behave thanks to this formulation:

- Perfect match between both pairs:

In order for both pairs of Diracs to coincide at the same time, a and b must satisfy the following conditions:

$$\begin{cases} \frac{-1+b}{a} = -1 \\ \frac{1+b}{a} = 1 \end{cases} \Leftrightarrow \begin{cases} b = 0 \\ a = 1 \end{cases} \tag{20}$$

- Match of one pair:

For only one Dirac of ∇u_0 to coincide with one Dirac of ∇v_0 , a and b need to satisfy one of the following conditions:

$$\begin{aligned}
\frac{1+b}{a} = 1 &\Leftrightarrow a - b = 1 \\
\frac{-1+b}{a} = -1 &\Leftrightarrow a + b = 1 \\
\frac{1+b}{a} = -1 &\Leftrightarrow a + b = -1 \\
\frac{-1+b}{a} = 1 &\Leftrightarrow b - a = 1
\end{aligned} \tag{21}$$

- In any other case, both pairs are separate, which leads to a null functional.

This heuristic study shows three configurations:

- One unique case ($a = 1$ et $b = 0$) for which both pairs of Diracs masses are perfectly aligned, meaning that both edges are correctly registered. Intuitively, this is when the functional is at its maximum, although it is not possible to formally evaluate its value due to the product of Diracs.
- 4 linear relations between a and b for which the signals have only one pair of edges out of two that matches. These relationships reflect an infinite number of local maxima, whose value is assumed to be lower than the perfect registration.
- No match between any edge, resulting in a null functional.

In order to describe mathematically the behavior of the functional and separate the global maximum from the local maxima, we study an approximation of the problem on differentiable signals that represents a differentiable approximation of the box function.

Theoretical registration

In order to deal with differentiable signals, we use the following approximation of the Heaviside function [2]:

$$H_\alpha(x) = \begin{cases} \frac{1}{2} \left(1 + \frac{x}{\alpha} + \frac{1}{\pi} \sin \frac{\pi x}{\alpha} \right) & \text{if } |x| \leq \alpha \\ 1 & \text{if } x > \alpha \\ 0 & \text{if } x < -\alpha \end{cases} \tag{22}$$

This function is differentiable, and its derivative is given by:

$$\delta_\alpha(x) = \begin{cases} \frac{1}{2\alpha} (1 + \cos \frac{\pi x}{\alpha}) & \text{if } |x| \leq \alpha \\ 0 & \text{if } |x| > \alpha \end{cases} \tag{23}$$

When $\alpha \rightarrow 0$, $\delta_\alpha \rightarrow \delta$ and $H_\alpha \rightarrow H$ where δ refers to a Dirac distribution and H the Heaviside function. In practice, α would tend to 0 to simulate a discrete edge.

Based on the previous definition of the reference signal v_0 , we can now rely on its continuous approximation:

$$v(x) = H_\alpha(x + 1) - H_\alpha(x - 1) \tag{24}$$

which is a box function on the interval $[-1 - \alpha; 1 + \alpha]$. Its derivative ∇v is given by:

$$\nabla v(x) = \delta_\alpha(x + 1) - \delta_\alpha(x - 1) \tag{25}$$

The signal to register and its derivative become:

$$\begin{aligned}
u_{a,b}(x) &= v(ax - b) \\
&= H_\alpha(ax - b + 1) - H_\alpha(ax - b - 1), \\
\nabla u_{a,b}(x) &= a (\delta_\alpha(ax - b + 1) - \delta_\alpha(ax - b - 1)).
\end{aligned}$$

Note that when $a > 1$, it results in a negative zoom, which means that the size of the support is reduced by a factor a , while the height of the peaks is multiplied by a factor a , hence heightened. On the contrary when $a < 1$ it is a positive zoom: the support is stretched by a factor $1/a$ and the height of the peaks is reduced. Figure 4 shows the effect of a zoom ($a > 1$ et $a < 1$) on the function u and its derivative.

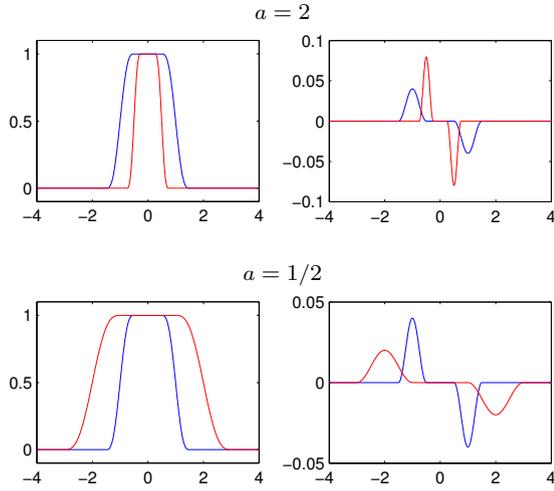


Fig. 4 Approximation with $\alpha = 0.5$ of the box function (on the left, in blue) and its derivative (on the right, in blue) and zoom with different values of a (in red).

The functional we seek to maximize becomes:

$$\begin{aligned}
 F(a, b) &= \int_{\mathbb{R}} |\nabla u_{a,b}(x) \cdot \nabla v(x)| \, dx \\
 &= a \int_{\mathbb{R}} (\delta_{\alpha}(x+1) + \delta_{\alpha}(x-1)) \\
 &\quad \cdot (\delta_{\alpha}(ax-b+1) + \delta_{\alpha}(ax-b-1)) \, dx \\
 &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1)\delta_{\alpha}(ax-b+1) \\
 &\quad + \delta_{\alpha}(x+1)\delta_{\alpha}(ax-b-1) \\
 &\quad + \delta_{\alpha}(x-1)\delta_{\alpha}(ax-b+1) \\
 &\quad + \delta_{\alpha}(x-1)\delta_{\alpha}(ax-b-1) \, dx \\
 &= F_1(a, b) + F_2(a, b) + F_3(a, b) + F_4(a, b)
 \end{aligned} \tag{26}$$

with:

$$\begin{aligned}
 F_1(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1)\delta_{\alpha}(ax-b+1) \, dx, \\
 F_2(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1)\delta_{\alpha}(ax-b-1) \, dx, \\
 F_3(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x-1)\delta_{\alpha}(ax-b+1) \, dx, \\
 F_4(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x-1)\delta_{\alpha}(ax-b-1) \, dx.
 \end{aligned} \tag{27}$$

Each of the sub-functionals F_1, \dots, F_4 can be studied separately in order to determine the conditions on a and b for the integrals to be maximal, and the close form of $F(a, b)$.

Proposition 1 *The functional F can be expressed under the following form:*

$$F(a, b) = \begin{cases} \frac{3}{2\alpha} & \text{if } a = 1 \text{ and } b = 0, \\ \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a} \right) & \text{if } a > 1 \text{ and } \begin{cases} a + b = 1 \\ \text{or } a + b = -1 \\ \text{or } a - b = 1 \\ \text{or } a - b = -1 \end{cases} \\ \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a \right) & \text{if } a < 1 \text{ and } \begin{cases} a + b = 1 \\ \text{or } a + b = -1 \\ \text{or } a - b = 1 \\ \text{or } a - b = -1 \end{cases} \\ 0 & \text{otherwise.} \end{cases} \tag{28}$$

Besides, it achieves its global maximum $\frac{3}{2\alpha}$ for $a = 1$ and $b = 0$.

A proof of this proposition is given in appendix B. This confirms the intuitive study conducted with the Dirac distributions in the first part, revealing a global maximum at the expected value $(a, b) = (1, 0)$ and linear subspaces of local maxima.

4.2 Assessment of the performance of the proposed criterion

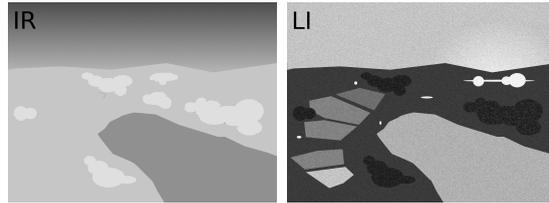


Fig. 5 Synthetic images of IR and LI modalities.

After showing that the proposed criterion is theoretically able to recover the optimal transformation parameters for the registration problem, this section aims at validating the proposed metric in practice. For the numerical evaluation, we have created synthetic images that reflect the characteristics of the involved modalities, displayed on Figure 5.

4.2.1 Study of the performances

First, we study on synthetic images the ability of the proposed criterion to find the optimal

Metric	Noiseless images	Poisson noise, PSNR \approx 18dB	Poisson noise, PSNR \approx 15dB
Cross-correlation [20]	73%	77%	77%
Mutual information [25]	6%	68%	61%
MI-G [19]	91%	100%	100%
MIND [10]	100%	52%	34%
Weighted cos [23]	100%	100%	100%
NGF [8]	100%	100%	95%
Proposed metric	100%	100%	100%

Table 1 Estimation of the percentage of correctly estimated transformations obtained on synthetic images with the cross-correlation (CC), the mutual information (MI), the combined edge-based/mutual information metric (MI-G), the MIND descriptor (MIND), the weighted cosine, the normalized gradient Field (NGF) and the proposed metric. The LI image is either noiseless or corrupted by a Poisson noise, so that its initial PSNR is around 18dB then 15dB.

transformation parameters. We simulate a hundred transformations with random translation and zoom parameters, distributed uniformly between -40 and +40 pixels for the horizontal and vertical translations, and between -0.4 and +0.4 for the zoom coefficient z , then we estimate the transformation using an exhaustive search on the 3-dimensional search parameter space. We round off the results according to the precision of the exhaustive search (two pixels for the translation and 0.05 for the zoom coefficient), then we count the number of correctly estimated transformations for each method, which gives a percentage of success. We also study its robustness to noise, since the optical images are corrupted by a strong non-Gaussian noise: we add Poisson noise to the LI images in order to reach a PSNR of about 18dB then 15dB, to reflect the natural degradations of this modality in night vision conditions. Then we compare the performances of our proposed edge-based metric to the classical methods described in section 1.2: the cross-correlation metric (CC) defined in equation (1), the mutual information (MI) computed from equation (2) the combined edge-based/mutual information (MI-G) based on equation (8), the MIND descriptors from (9), and the different edge-based metrics that have been described : the weighted cosine of equation (3), the Normalized Gradient Fields (NGF) from equation (6), and the proposed method.

This evaluates the ability of each metric to attain its global maximum with the optimal transformation. Tables 1 displays the percentage of correctly estimated transformations over the hundred randomly simulated ones with each metric, on clear then Poisson-corrupted images. The results show that the edge-based methods provide a more reliable estimation, that is also more robust to noise. As described in

the first section, the mutual information and the cross-correlation are shown not to be suited to the multi-modal characteristics of the IL/IR images. When using the combined MI-G metric [19], the edge information allows the registration to be performed accurately, but the mutual information is actually not beneficial here. Indeed, when the noise power increases, so does the minimum value of the image gradients, which increases the power of the edge information over the mutual information (based on equation (8)), which might explain why it performs more accurately on noisy images. The MIND descriptors are suited to the multimodality of the images, however the descriptors proposed in [10] are adapted to measure similarities of Gaussian-corrupted patches, as in the original NL-means. The three edge-based methods are the ones that are the more accurate here, even though the NGF metric does fail in the presence of strong noise. This can be explained by the normalization of the gradients that makes it more sensitive to noise, and the selection of the noise parameter ϵ that might not be effective enough to distinguish edges from strong noise.

4.2.2 Study of the stability

Since we perform an optimization scheme for the search for the optimal parameters, the stability of the metric, i.e. its sensitivity to local maxima and its behavior around the global maximum, is a crucial point. Indeed, for the gradient ascent to be liable, the energy needs to be as concave as possible, smooth, and it must not suffer from too many local maxima that would lure the gradient ascent away from the optimum.

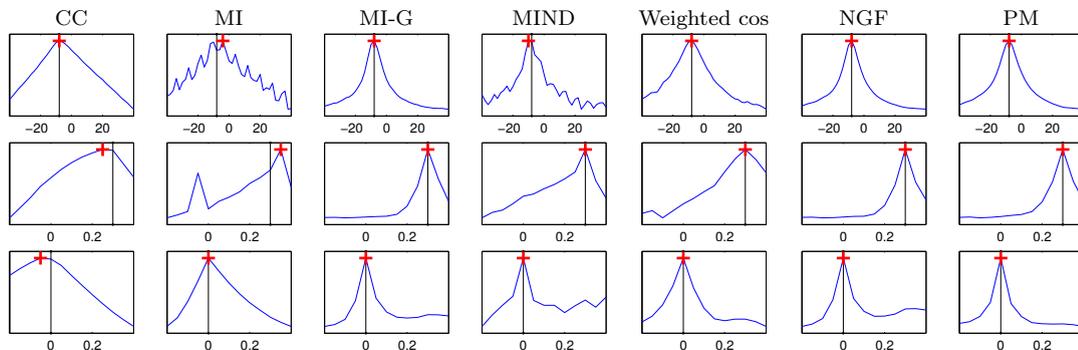


Fig. 6 Evolution of the metrics as a function of one parameter, the other two being fixed to the optimal value. First line: noiseless images, optimal parameters $[t_1, t_2, z] = [8, -8, 0.35]$, search for the horizontal translation. Second line: Poisson-corrupted images (PSNR ≈ 15 dB), optimal parameters $[t_1, t_2, z] = [24, -32, 0.3]$, search for the zoom parameter. Third line: Poisson-corrupted images (PSNR ≈ 15 dB), optimal parameters $[t_1, t_2, z] = [-4, 2, 0]$, search for the zoom parameter. From left to right: Cross-correlation [20], Mutual information [25], combined Edge-based/Mutual information [19], MIND [10], Weighted cosine [23], Normalized Gradient Fields [8] and proposed metric (11).

The next experiment consists in evaluating the behavior of the different metrics regarding the variation of one parameter. We simulate a known transformation on the synthetic images, and we compute the criterion while testing only one parameter, the other two being fixed to the correct value. Figure 6 displays the evolution of the metrics in the case of the cross-correlation metric, the mutual information metric, the combined edge-based mutual-information metric [19], the MIND descriptors [10], the weighted cosine metric [23], the normalized Gradient Fields [8] and the proposed metric. On the first line, the LI image was corrupted by Poisson noise (initial PSNR ≈ 15 dB) with optimal parameters $[t_1, t_2, z] = [8, -8, 0.35]$, and the tested parameter was the horizontal translation parameter t_2 . On the second line, the experiment was conducted on noiseless images with initial parameters $[t_1, t_2, z] = [24, -32, 0.3]$ and the tested parameter was the zoom coefficient z . On the bottom line, the LI image was also corrupted by Poisson noise (initial PSNR ≈ 15 dB) with initial parameters $[t_1, t_2, z] = [-4, 2, 0]$, and the tested parameter was the zoom coefficient z . These figures illustrate the ability of each metric to reach its maximum at the optimal transformation but also its behavior around the maximum. The cross-correlation does offer smooth variations but it tends to misplace the optimum, in the second and third cases for example. The mutual information is shown to be unstable and suffers from multiple local maxima, which can introduce errors (first and second lines). The contribution of the edge in-

formation does stabilize the metric with combined MI-G, but the solely edge-based metrics are equally efficient and simpler to optimize (pixel wise metric rather than histogram-based). These plots also confirm the inability for the MIND descriptors to deal with high level of non-Gaussian noise. Finally, the three edge-based metrics (weighted cosine, NGF and the proposed method) show reliable plots, but the weighted cosine and the NGF metrics are slightly less stable when further from the maximum, as shown on the first and third transformations.

Both studies on performance and stability lead us to confirm the theoretical results and to validate the proposed metric.

4.3 Study of the optimization scheme

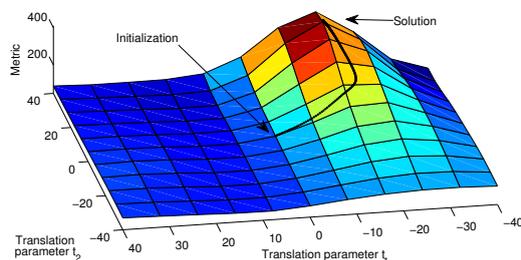


Fig. 7 Similarity map for a range of translation parameters, at a fixed zoom, and estimated parameters at each iteration of the gradient ascent. We can see that the optimization scheme does maximize the metric, and that it converges towards the optimum.

The above theoretical and experimental studies have sought to validate the proposed metric in terms of performance, robustness and relevance. We have shown that the maximization of the proposed metric does lead to the optimal transformation, both in theory and in practice. We have also studied its robustness to noise and its behavior regarding the evolution of one transformation parameter, and we have demonstrated its ability to find the optimal transformation when searching the whole parameter space.

Now that the performance of the functional has been validated, we focus on the optimization scheme that we have developed. Indeed, even though the functional is not concave, we can study its ability to attain the global maximum within a gradient ascent scheme, provided that the initialization is close enough to the solution. In our night vision context, we benefit from a video flux that ensures that the registration performed for the previous frames is a good guess for the next couple of images.

Figure 7 displays the map of the metric computed for a fixed zoom parameter. The black line shows the path of the estimated translation parameters t_1 and t_2 at each iteration of the gradient ascent scheme. We can see that each step does maximize the metric, and that the algorithm converges towards the optimum. Besides, even though the criterion is not strictly concave, this map shows that around the optimum the metric behaves well, which guarantees that with a “good” initialization the gradient ascent will converge to the global maximum.

Such a “good” initialization is actually guaranteed by the fact that we treat videos, so we can benefit from the estimation of the frame before as described in section 3.2. Based on the synthetic images, we have simulated movement on the IR image with the translation parameters t_1 and t_2 between -40 and +40 pixels, and the zoom parameter z between 0 and 0.4, evolving with time. Then the registration with the LI image was performed using the gradient ascent for each frame. We start the first frame with a coarse exhaustive search, and we use the estimated parameters for each frame as an initialization for the next one. This guarantees that the initialization is not too far from the optimum, and it allows to converge in a limited number of iterations. Figure 8 shows the evolution of each simulated parameter with time

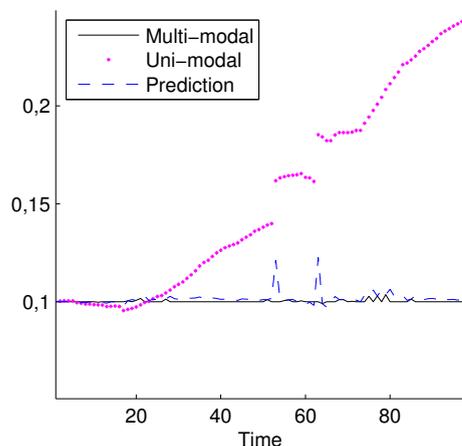


Fig. 10 Evolution of the estimated zoom parameter for a sequence of real images using the multi-modal registration method (black line), the sole propagation of the uni-modal information (purple stars), and initialization before the gradient ascent (blue dashes).

(blue line), and the red stars show the estimation of these parameters for each frame. The algorithm performs well on all the sequence. In fact, the average error on the estimation is less than 1 pixel for the translation and 0.003 for the zoom. These experiments show that maximizing the proposed metric is relevant since it does lead to the optimum transformation, and that the proposed gradient ascent scheme is successful.

4.4 Study of the temporal scheme

If the relative position between the two cameras evolves more suddenly, the transformation parameter will vary abruptly and the gradient ascent might not follow this fast change. This is the reason why we have introduced the temporal validation scheme that thanks to the posture detection, evaluates the movement of each camera independently and predicts the registration parameter for the next frames, or propagates the movements. If the movement is global and can be assumed to be a movement of the scene, for example due to a gust of wind, then the relative position between the two cameras is unchanged and the registration performance is unaltered. If the change applies only to one camera, for example due to a sudden movement of the pilot’s head, then the registration parameters will be altered and the sole gradient ascent might not catch such variations, hence the point of the temporal validation scheme. Figure 9 displays the evolution (blue line) and es-

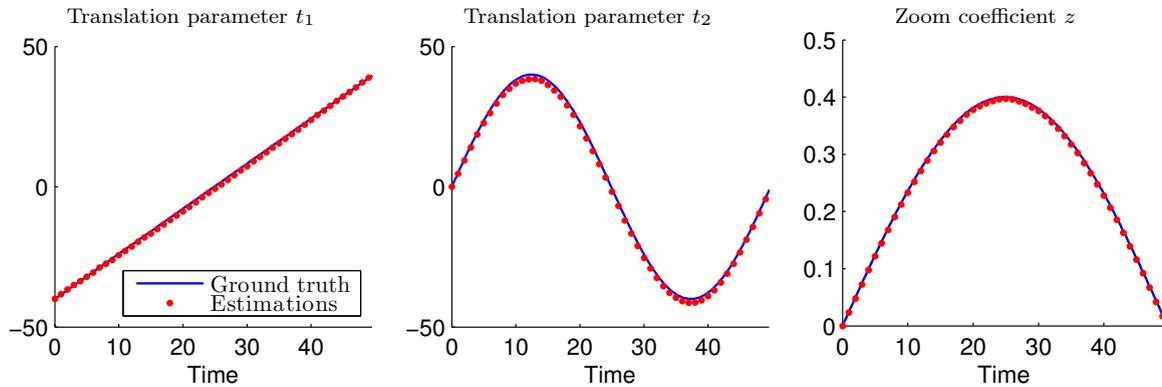


Fig. 8 Evolution of each parameter (t_1, t_2, z) of the transformation during the sequence and estimation (in red stars) computed using the gradient ascent. Some of the errors are due to the fact the evolution of the parameters is continuous, so the values are not integers, which generates approximations in the transformation.

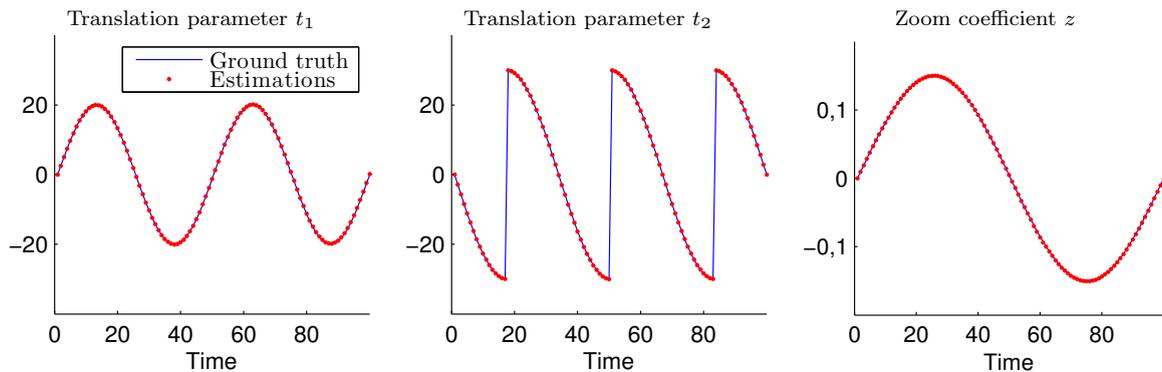


Fig. 9 Evolution of each parameter (t_1, t_2, z) of the transformation during the sequence and estimation (in red stars) computed using the gradient ascent. The discontinuities observed with the horizontal translation parameter t_2 simulate sudden movements of the head, that are properly handled thanks to the temporal validation scheme.

timization (red stars) of each simulated parameter, when the horizontal translation parameter suffers from discontinuities that simulate sudden movements of the pilot’s head. This shows that even when confronted to discontinuities, the proposed algorithm is able to register accurately the two cameras, thanks to the uni-modal based prediction.

Besides, both uni-modal and multi-modal estimations, ie the multi-modal estimation T_t between the two cameras and the uni-modal estimations A_t and B_t of the movement of each camera separately (see figure 3), complete each other. Indeed, in our temporal framework we have introduced the uni-modal estimation as a safety net that checks that the gradient ascent has worked properly, but it is also true the other way around: the accumulation of consecutive uni-modal estimates introduces a drift in the estimated transformation and the multi-modal step is thus mandatory to adjust the estimation. Figure 10 displays the evolution of the

zoom parameter for a sequence of (real) images. The black line displays the estimated parameter with our multi-modal method, the purple stars show the estimation using only the uni-modal information by propagating the initialization, and the dashed blue line is the initialization parameter from \tilde{T}_t at each step. This plot shows two things: first, that the sole uni-modal information introduces a drift, so the uni-modal registration is inaccurate. It shows also that at (nearly) each step the multi-modal step has a correction to do compared to the estimation \tilde{T}_t (see the moderate difference between the black and blue lines).

This justifies the combination of both uni-modal and multi-modal methods to guarantee precision and liability.

5 Projective model

When the helicopter flies at high altitude, the assumption that the transformation between

the two modalities can be modeled by a uniform zoom and a translation (or more generally by an affine transformation) can be verified. However, when the helicopter flies at lower altitude, the perspective is different between both cameras, so a projective model has to be adopted.

5.1 Projective geometry

A projective transformation [9] is described by the homography matrix

$$H = \begin{pmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & 1 \end{pmatrix} \quad (29)$$

that has 8 degrees of freedom.

The equation of the transformation is:

$$\begin{pmatrix} wx' \\ wy' \\ w \end{pmatrix} = H \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \begin{pmatrix} h_{11}x + h_{12}y + h_{13} \\ h_{21}x + h_{22}y + h_{23} \\ h_{31}x + h_{32}y + 1 \end{pmatrix}, \quad (30)$$

then we revert to x' and y' by normalizing by w :

$$\begin{cases} x' = \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + 1} \\ y' = \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + 1} \end{cases} \quad (31)$$

We can simplify the expression using only the first two coordinates:

$$X' = HX = \begin{pmatrix} \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + 1} \\ \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + 1} \end{pmatrix} \quad (32)$$

This type of transformation is a generalization of the affine model, and it includes the transformation model we considered until then, but also the rotations and the changes of perspective.

Even though it is possible to restrain the space of the sought parameters, an exhaustive search would still require a higher number of dimensions which makes it computationally difficult in real-time. The gradient ascent is then ever more important in the extended problem.

5.2 Gradient ascent

The functional $F = F(H)$ that we seek to maximize now relies on 8 parameters, so we have to compute 8 partial gradients regarding each

parameter. We note H_{11} the variation on parameter h_{11} :

$$H_{11} = \begin{pmatrix} h_{11} + \alpha & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{pmatrix} \quad (33)$$

so

$$\begin{aligned} H_{11}X &= \begin{pmatrix} \frac{(h_{11} + \alpha)x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + h_{33}} \\ \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + h_{33}} \end{pmatrix} \\ &= HX + \alpha \begin{pmatrix} \frac{x}{h_{31}x + h_{32}y + h_{33}} \\ 0 \end{pmatrix} \end{aligned} \quad (34)$$

We can write in a similar way the variations H_{12}, \dots, H_{23} .

For the parameters that intervene in the denominator, we need to perform a linearization and we have:

$$H_{31}X = \begin{pmatrix} \frac{h_{11}x + h_{12}y + h_{13}}{(h_{31} + \alpha)x + h_{32}y + h_{33}} \\ \frac{h_{21}x + h_{22}y + h_{23}}{(h_{31} + \alpha)x + h_{32}y + h_{33}} \end{pmatrix}, \quad (35)$$

which leads to the following result:

$$\begin{aligned} H_{31}X &= HX - \alpha \frac{x}{h_{31}x + h_{32}y + h_{33}} HX + o(\alpha) \\ &= HX \left(1 - \alpha \frac{x}{h_{31}x + h_{32}y + h_{33}} \right) + o(\alpha) \end{aligned}$$

And also:

$$\begin{aligned} H_{32}X &= \begin{pmatrix} \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + (h_{32} + \alpha)y + h_{33}} \\ \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + (h_{32} + \alpha)y + h_{33}} \end{pmatrix} \\ &= HX - \alpha \frac{y}{h_{31}x + h_{32}y + h_{33}} HX + o(\alpha). \end{aligned}$$

We can then re-inject each of this variation calculation in the computation of the metric, leading to a 8-dimensional gradient.

Projective transformations are very unstable due to the non-linearity of the denominator, which makes a direct estimation quite difficult. Inspired by [15], we initialize the first couple of frames with a non-projective transformation, which offers more stability and allows to perform an exhaustive search on a 3-dimensional space instead of 8, then the projective dimension is included in the gradient ascent for the next frames.

	Exhaustive search			Gradient ascent (Zoom/Translation)			Gradient ascent (Projective)		
	Matlab	C++	GPU	Matlab	C++	GPU	Matlab	C++	GPU
Iteration time (ms)									
1600 × 1200	756.37	153.15	7.46	1455.71	195.01	9.38	1657.97	255.23	12.59
800 × 600	175.53	37.26	2.99	328.52	44.56	3.28	394.28	56.36	4.91
400 × 300	46.5	9.19	0.998	88.36	11.38	1.22	99.05	15.09	1.66
Total registration time (s)									
	≈ 10000 iterations			≈ 100 iterations			≈ 100 iterations		
1600 × 1200	7563.7	1531.5	74.6	145.571	19.501	0.938	165.797	25.523	1.259
800 × 600	1755.3	372.6	29.9	32.852	4.456	0.328	39.428	5.636	0.491
400 × 300	465	91.9	9.98	8.836	1.138	0.122	9.905	1.509	0.166
Frame rate (Hz)									
1600 × 1200	0.0001	0.0007	0.0134	0.0069	0.0513	1.0661	0.0060	0.0392	0.7943
800 × 600	0.0006	0.0027	0.0334	0.0304	0.2244	3.0488	0.0254	0.1774	2.0367
400 × 300	0.0022	0.0109	0.1002	0.1132	0.8787	8.1967	0.1010	0.6627	6.0241

Table 2 Registration time computed for the exhaustive search, the standard gradient ascent scheme (dealing with only a zoom and a translation), and the projective gradient ascent scheme, depending on the size of the image and the implementation. A GPU implementation of the gradient ascent scheme allows to perform registration in less than a second.

6 Results

6.1 Computational time and implementation

The proposed method has been implemented first on Matlab, then in C++ and optimized with Graphics Processing Unit (GPU) to accelerate the registration time. The initial LI images are of size 1600 × 1200 pixels and the IR images are of size 768 × 576. First we resize the images to the same dimensions (using bicubic interpolation), then we perform a down-sampling in order to reduce the size of the images, hence the computational time.

To find the solution with an exhaustive search, we need to compute the value of the metric for each tested set of parameters, which implies applying the associated transformation to the current image, then computing the metric (which involves calculating the gradient of each image). This step is repeated for every set of parameters of the search space, that includes at least 10000 possibilities (in the non-projective case)!

For a gradient ascent scheme, for each iteration step the registered image is computed in order to evaluate the gradient of the metric, which also involves computing the image gradients. Experiments have shown that 100 iterations allow the gradient ascent to converge, and this number can even be reduced when associated to a temporal scheme as in section 3.2 where the initialization is refined by a monomodal registration.

The gradient ascent scheme is all the more interesting compared to the efficiency of the exhaustive search when the number of transformation parameters to estimate becomes important, for example in the projective case. In order to illustrate the computational complexity involved with each method, we display in table 2 the computational time as a function of the image size (that depends on the down-sampling factor) needed for one step of the computation of the solution: either one iteration of the gradient ascent ascent or one computation of the metric for one set of parameters. Then by taking into account the average number of iterations needed (number of iteration steps for the gradient ascent scheme or size of the parameters search space), it gives an indication of the registration time for one image, depending on the size. The computational time can then expressed in terms of frame rate, ie the number of frames that can be processed in a second, that is shown to be up to 8 frames per second.

6.2 Experimental validation on real experimental conditions

In order to evaluate the different metrics on real images without knowing the optimal transformation, we have developed a method based on manually selected points. We select on the reference IR image and the LI image to register ten pairs of characteristic points (the whole difficulty being to find reference points that can be identified in both images), and we measure

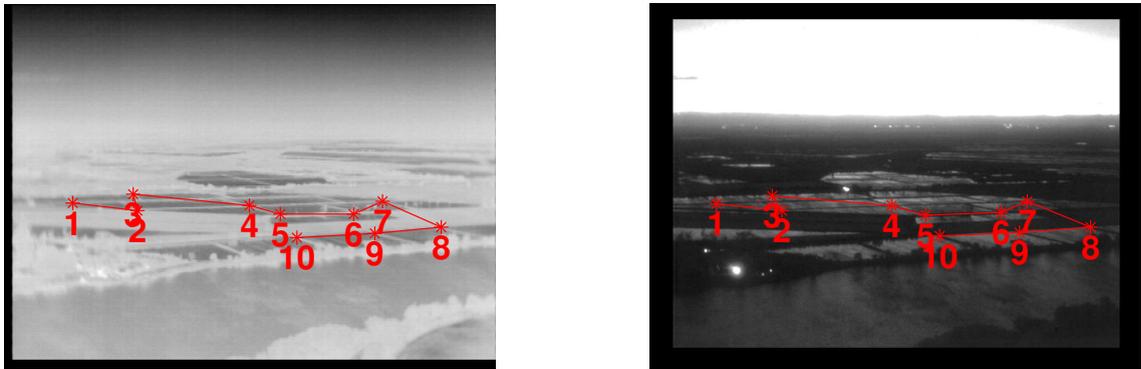


Fig. 11 Example of control points selected on the reference IR image, and associated points localized on the registered image.

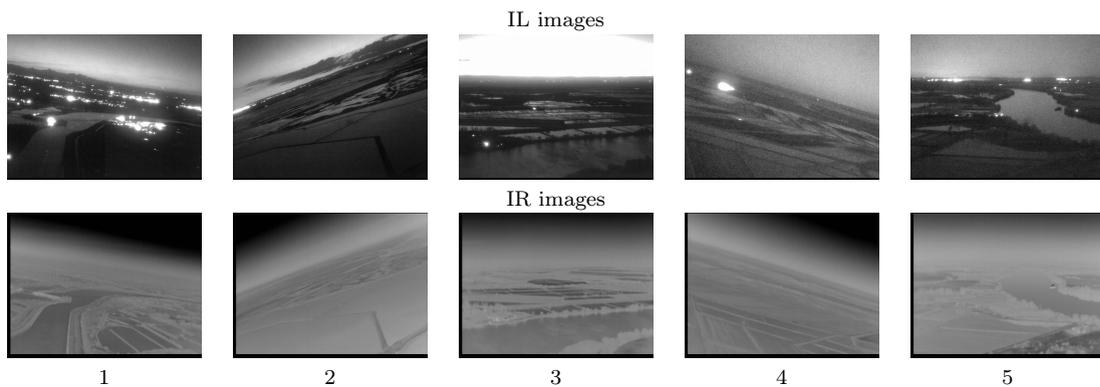


Fig. 12 Extract of the image sequences that have been used for the landmarks tests. According to the conditions and level of light, the images reflect different characteristics or artifacts.

Image series	Image to register	CC [20]	MI [25]	MI-G [19]	MIND [10]	Weighted cos [23]	NGF [8]	Proposed metric
1	31.53	33.36	13.34	17.11	11.75	6.69	40.42	6.65
2	31.23	40.76	11.80	13.16	11.81	11.96	11.81	11.26
3	19.79	28.64	7.97	5.86	4.25	4.04	12.13	4.42
4	25.57	31.99	23.93	10.23	8.62	8.85	13.13	8.58
5	29.33	35.02	10.69	9.01	9.40	8.10	9.44	7.80

Table 3 Average distance between the pixel coordinates of the control points from the reference image to the image to register and to the registered image, using either the cross-correlation metric (CC), the mutual information (MI), the combined edge-based/mutual information (MI-G), the MIND descriptors, the weighted cosine, NGF or the proposed metric. The best average errors are displayed in bold, while the italic indicates results that visually performed a satisfying registration.

the distance (in pixels) between the pixel coordinates for each pair of points. Then we measure the distance between the pixel coordinates from the reference image and those from the registered image. We repeat this measurement for a sequence of ten images for each series, then we average the distances to produce an average registration error (in pixels). Figure 11 displays an example of these landmarks manually selected on the reference IR image, and the corresponding points in the registered image. We have selected five series of images displayed on figure 12, that reflect different prop-

erties of the night vision. The first three image sequences suffer from a relatively small amount of noise and offer good contrast, while the image sequences 4 and 5 are strongly affected by Poisson noise and suffer from a lower contrast. The first and fourth image sequence are also affected by light sources artifacts that disrupt the dynamic and conceal some information. On the other hand, the IR images are poorly contrasted and might offer little information, as for the sequences 2 and 4. This experiment has been conducted using the cross-correlation (1), the mutual information (2), the combined edge-based

and mutual information (8), the MIND descriptors (9), the weighted cosine metric (3), NGF (6) and our proposed method. The average error distance between the pixels coordinates of the reference image and the registered ones is displayed in table 3. Our proposed edge-method is shown to perform a more accurate registration in most cases. The edge-based metric based on the weighted cosine (3) and the MIND descriptor (9) can also give accurate registration, but they seem respectively more sensitive to low contrast or to the presence of noise and artifacts.

6.3 Registration on real data

Figure 13 displays an infra-red image, an optical image and the registered optical image in the IR coordinates issued from three different sequences. For each sequence, the images have been resized to the common resolution of 1024×768 pixels, then down-sampled by a factor 2. The parameters for the gradient ascent have been initialized using a coarse exhaustive search, then the gradient steps are set to $1e-4$ for the translation parameters and $1e-8$ for the zoom parameters. These steps have been manually optimized, but they are fixed for all the data we have tested. Besides, in order to simulate the estimation of the temporal movement of each camera that should be provided by the posture detection and inertial systems, we have simulated this estimation using Motion 2D [17], which is a fast and simple method to perform registration of a single modality through time.

In order to illustrate the accuracy of the registration, we have performed an edge detection on the IR image, and we have printed these edges on the optical images, to check that the edges are correctly aligned. We have also created some hybrid mosaic images that display alternatively stripes of the reference IR image and stripes of the LI image, before then after registration. These mosaics allow at a single glance to assess the accuracy of the registration, by observing the continuity of the edges from one modality to the other. The complete registered sequences are available for download¹.

¹ <http://image.math.u-bordeaux1.fr/Registration>

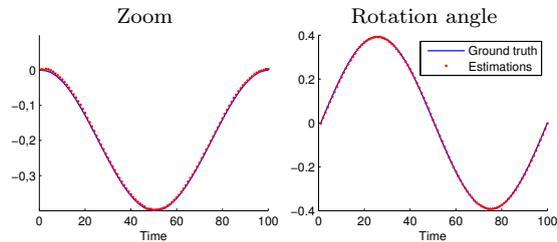


Fig. 14 Simulation of a rotating LI image sequence with varying zoom and rotation angle parameters (blue line) and estimation for the registration (red stars).

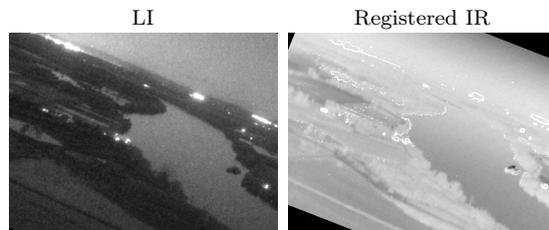


Fig. 15 Example of a LI image transformed with a zoom and a rotation, and the associated registered IR image. The edges of the LI image have been superimposed onto the IR one in order to illustrate the accuracy of the registration.

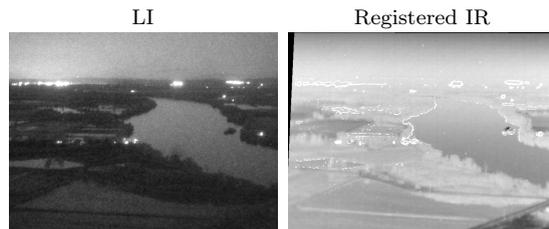


Fig. 16 Example of a LI image transformed with a zoom, a translation in both directions and a horizontal projection, and the associated registered IR image. The edges of the LI image have been superimposed onto the IR one in order to illustrate the accuracy of the registration.

6.4 Projective simulations

We have simulated transformations on already registered (real) images. Since the projective geometry includes rotations and changes of perspective, we have simulated such movements on a LI image sequence, then performed the registration of the IR image for each frame. The knowledge of the true transformation parameters allows to check that the estimation is accurate.

Figure 14 displays the simulated and estimated zoom and rotation angle parameters during the sequence. Figures 15 and 16 show the transformed LI image subjected to either a zoom and a rotation or a zoom, translation and

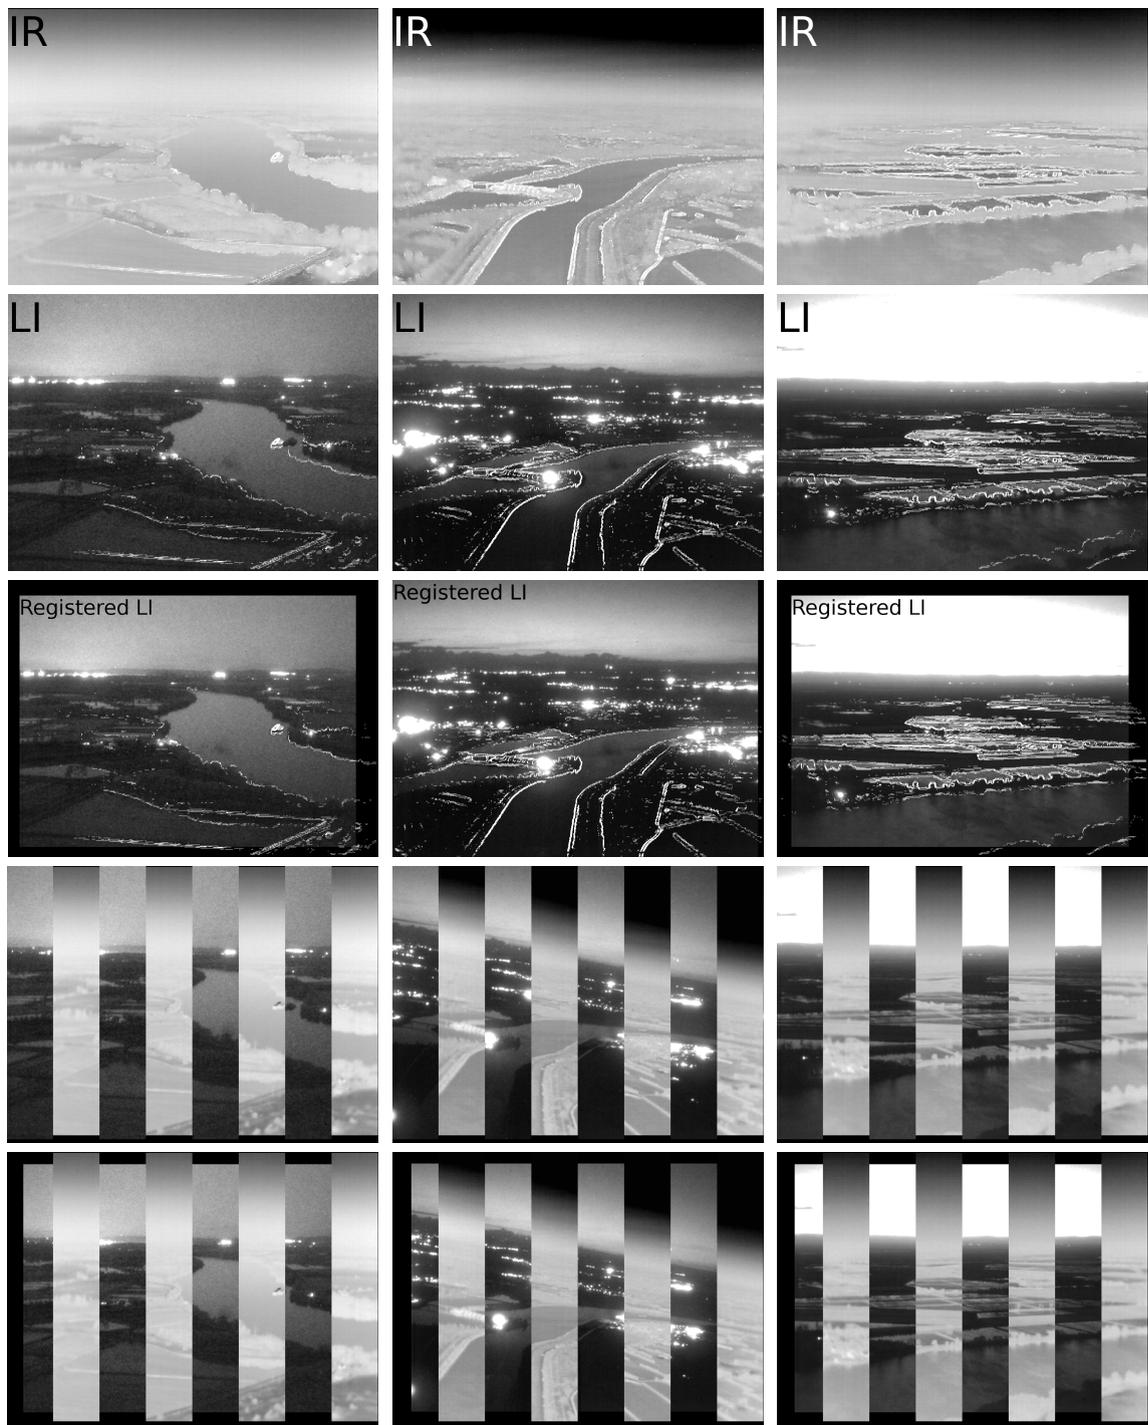


Fig. 13 Results obtained on a real data sets using the proposed metric. The reference IR image is displayed on the first line, the current LI image to register on the second line, and the registered LI image on the third line. Edges extracted from the IR images are superimposed on each image for visual assessment of the quality of the registration. On the fourth line is displayed a mosaic image that displayed alternatively on each band either the IR or the LI image to register, to better assess the discrepancy between the two modalities, while the fifth line displays the Mosaic image computed with the registered LI image to judge the alignment of the edges. The complete registered sequences are available at <http://image.math.u-bordeaux1.fr/Registration>.

horizontal projection, and the corresponding IR registered images.

These figures illustrate that the general projective model can accurately encompass many

kinds of transformations that are likely to be encountered in-flight.

Nonetheless, the approach that we have developed here assumes the scene to be planar,

so a parametric transformation can be estimated. If the helicopter flies at very low altitude for example, the presence of buildings or trees might violate this assumption. In this case, the maximum of the functional corresponds to the matching of the biggest possible number of edges, which would consist in aligning the major plan of the scene, hence treating the rest as outliers. Such configurations might be investigated further in future studies.

7 Conclusion and perspectives

In this paper, we have presented a new multi-sensor registration method based on the edge alignment principle. We have developed a new algorithm that aligns the edges that appear in both modalities by performing a gradient ascent scheme that provides a fast algorithm. Coupled with a temporal implementation that ensures more liability and provides error control, our proposed method is shown to be robust and fast compared to a standard exhaustive search, and the algorithm can proceed up to 8 frames per second. Theoretical and experimental studies show that the criterion is suitable and liable, and results on real data validate the night vision application.

Besides, the computational performance could be further improved by exploiting the temporal scheme: by studying the behavior of the uni-modal registration, it would be possible to perform the multi-modal registration every n frames (depending on the observed drift), using only the uni-modal information in-between, which would offer a major gain of time that could be interesting for our operational applications. This might be studied in future work.

A Computation of the gradient of Functional (15)

We focus on the continuous form of the functional:

$$F(T) = \int_{\Omega} |\nabla u(T(X)).\nabla v(X)| \, dX,$$

where T is the transformation we seek to optimize.

If we define a small displacement S , we have: $F(T+S) = \int_{\Omega} |\nabla u(T(X)+S(X)).\nabla v(X)| \, dX$ and: $\nabla u(TX+SX) = \nabla u(TX) + D^2u(TX)(SX) + o(S)$.

We have:

$$F(T+S) = \int_{\Omega} |\nabla u(T(X)+S(X)).\nabla v(X)| \, dX \quad (36)$$

Using the first order expansion, we have:

$$\begin{aligned} & \int_{\Omega} |\nabla u(T(X)).\nabla v(X) \\ & + D^2u(TX)(SX).\nabla v(X)| \, dX \\ & = \int_{\Omega} |\nabla u(T(X)).\nabla v(X)| \\ & \times \left| 1 + \frac{D^2u(TX)(SX).\nabla v(X)}{\nabla u(T(X)).\nabla v(X)} \right| \, dX \\ & = \int_{\Omega} |\nabla u(T(X)).\nabla v(X)| \\ & \times \left(1 + \frac{D^2u(TX)(SX).\nabla v(X)}{\nabla u(T(X)).\nabla v(X)} \right) \, dX \quad (37) \end{aligned}$$

so that

$$\begin{aligned} F(T+S) & = F(T) \\ & + \int_{\Omega} \sigma D^2u(TX)(SX).\nabla v(X) \, dX + o(S). \quad (38) \end{aligned}$$

with $\sigma = \text{sign}(\nabla u(T(X)).\nabla v(X))$. We now focus on a variation on each parameter t_1, t_2, z and we see the functional F as a function of each parameter:

$$F(t_1 + \alpha, t_2, z) - F(t_1, t_2, z) = \alpha \partial_1 F(t_1, t_2, z)$$

and we denote by T_{α} the perturbation on T of α on the first parameter t_1 , ie:

$$T_{\alpha} = \begin{pmatrix} 1+z & 0 & (t_1 + \alpha) \\ 0 & 1+z & t_2 \\ 0 & 0 & 1 \end{pmatrix}, \text{ and we have:}$$

$$T_{\alpha}X = TX + \alpha \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

Hence,

$$\nabla u(T_{\alpha}X) = \nabla u(TX) + \alpha D^2u(TX) \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} +$$

$o(\alpha)$ and:

$$\begin{aligned} & F(t_1 + \alpha, t_2, z) - F(t_1, t_2, z) \\ & = \alpha \int_{\Omega} \sigma D^2u(TX) \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} .\nabla v(X) + o(\alpha) \quad (39) \end{aligned}$$

and we deduce that:

$$\partial_1 F(t_1, t_2, z) = \int_{\Omega} \sigma D^2u(TX) \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} .\nabla v(X) \quad (40)$$

We obtain an analogous result for the second parameter t_2 :

$$\partial_2 F(t_1, t_2, z) = \int_{\Omega} \sigma D^2 u(TX) \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \cdot \nabla v(X) \quad (41)$$

For the zoom parameter z , we consider:

$$T_{\gamma} = \begin{pmatrix} 1+z+\gamma & 0 & t_1 \\ 0 & 1+z+\gamma & t_2 \\ 0 & 0 & 1 \end{pmatrix}, \text{ so:}$$

$$T_{\gamma} X = TX + \gamma \begin{pmatrix} x \\ y \\ 0 \end{pmatrix}$$

Hence,

$$\nabla u(T_{\gamma} X) = \nabla u(TX) + \gamma D^2 u(TX) \begin{pmatrix} x \\ y \\ 0 \end{pmatrix} +$$

$o(\gamma)$

and:

$$\begin{aligned} & F(t_1, t_2, z + \gamma) - F(t_1, t_2, z) \\ &= \gamma \int_{\Omega} \sigma D^2 u(TX) \begin{pmatrix} x \\ y \\ 0 \end{pmatrix} \cdot \nabla v(X) + o(\gamma) \end{aligned} \quad (42)$$

So we have:

$$\partial_3 F(t_1, t_2, z) = \int_{\Omega} \sigma D^2 u(TX) \begin{pmatrix} x \\ y \\ 0 \end{pmatrix} \cdot \nabla v(X) \quad (43)$$

Putting those three differentials together, we obtain an explicit form for:

$$\nabla F(t_1, t_2, z) = \begin{pmatrix} \partial_1 F(t_1, t_2, z) \\ \partial_2 F(t_1, t_2, z) \\ \partial_3 F(t_1, t_2, z) \end{pmatrix}. \quad (44)$$

B Proof of proposition 1: Theoretical analysis of the criterion

We have expressed the functional $F(a, b)$ that we seek to maximize as the sum of four sub-functionals:

$$\begin{aligned} F_1(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1) \delta_{\alpha}(ax-b+1) dx, \\ F_2(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1) \delta_{\alpha}(ax-b-1) dx, \\ F_3(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x-1) \delta_{\alpha}(ax-b+1) dx, \\ F_4(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x-1) \delta_{\alpha}(ax-b-1) dx. \end{aligned} \quad (45)$$

We can study each sub-functional separately in order to determine the conditions on a and b for each of them to be maximal.

Note that the parameter α that represents the width of the peaks is meant to tend to 0. The peaks issued from the derivative of the reference signal v are located at $+1$ and -1 , and their support is $[\pm 1 - \alpha, \pm 1 + \alpha]$. For the transformed signal $u_{a,b}$, the peaks are located in $\frac{\pm 1+b}{a}$ and the support is $[\frac{\pm 1+b-\alpha}{a}, \frac{\pm 1+b+\alpha}{a}]$, of half-width α/a .

Hence, when α tends to 0, the width of each peak tends to 0 (for the transformed signal, this implies that $a > 0$, which is relevant in practice).

This remark simplifies the problem: we can consider that if the centers of the peaks are not perfectly aligned, then it is possible to consider a small enough α such that the supports are disjoint. Hence, we split the proof into five steps and the study of the 4 sub-functionals is limited to the conditions on a and b for the centers of the peaks to be aligned.

Step 1:

$$F_1(a, b) = a \int_{\mathbb{R}} \delta_{\alpha}(x+1) \delta_{\alpha}(ax-b+1) dx$$

The support of $\delta_{\alpha}(x+1)$ is $[-1-\alpha; -1+\alpha]$, centered in -1 , and the support of $\delta_{\alpha}(ax-b+1)$ is $[\frac{-1+b-\alpha}{a}; \frac{-1+b+\alpha}{a}]$, centered in $-1+b/a$.

For the function F_1 to be non-null, we solve:

$$\frac{-1+b}{a} = -1 \Leftrightarrow a+b=1 \quad (46)$$

Besides, when condition (46): $a+b=1$ is fulfilled, F_1 can be expressed in closed form:

- $a > 1$:

If $a > 1$, the half-size of the support of $u_{a,b}$ is $\alpha/a < \alpha$, so the computation of F_1 is restricted to the interval $[-\alpha/a, \alpha/a]$, and we have:

$$\begin{aligned} F_1(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1) \delta_{\alpha}(ax-b+1) dx, \\ &= \frac{a}{4\alpha^2} \int_{-\frac{\alpha}{a}}^{\frac{\alpha}{a}} \left(1 + \cos \frac{\pi x}{\alpha}\right) \cdot \left(1 + \cos \frac{\pi ax}{\alpha}\right) dx, \\ &= \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a}\right). \end{aligned} \quad (47)$$

- $a < 1$:

If $a < 1$, the half-size of the support of $u_{a,b}$

is $\alpha/a > \alpha$, so the computation of F_1 is restricted to the interval $[-\alpha, \alpha]$, and we have:

$$\begin{aligned} F_1(a, b) &= a \int_{\mathbb{R}} \delta_{\alpha}(x+1) \delta_{\alpha}(ax-b+1) dx, \\ &= \frac{a}{4\alpha^2} \int_{-\alpha}^{\alpha} \left(1 + \cos \frac{\pi x}{\alpha}\right) \cdot \left(1 + \cos \frac{\pi ax}{\alpha}\right) dx, \\ &= \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a \right). \end{aligned} \quad (48)$$

- $a = 1$:

If $a = 1$, the condition (46): $a+b = 1$ implies that $b = 0$, so $u = v$, and we have:

$$F_1(1, 0) = \int_{-\alpha}^{\alpha} \left[\frac{1}{2\alpha} \left(1 + \cos \frac{\pi x}{\alpha}\right) \right]^2 dx = \frac{3}{4\alpha} \quad (49)$$

Conclusion :

$$F_1(a, b) = \begin{cases} \frac{3}{4\alpha} & \text{if } a = 1 \text{ and } b = 0 \\ \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a}\right) & \text{if } a+b = 1 \text{ and } a > 1 \\ \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a\right) & \text{if } a+b = 1 \text{ and } a < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (50)$$

Step 2:

$$F_2(a, b) = a \int_{\mathbb{R}} \delta_{\alpha}(x+1) \delta_{\alpha}(ax-b-1) dx$$

An analogous study on the support of $\delta_{\alpha}(x+1)$ and $\delta_{\alpha}(ax-b-1)$ leads to solving the following conditions on a and b for the support to intersect:

$$\frac{1+b}{a} = -1 \Leftrightarrow a+b = -1 \quad (51)$$

When condition (51): $a+b = -1$ is satisfied, F_2 can be computed in a similar way to F_1 :

$$F_2(a, b) = \begin{cases} \frac{3}{4\alpha} & \text{if } a = 1 \text{ and } b = -2 \\ \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a}\right) & \text{if } a+b = -1 \text{ and } a > 1 \\ \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a\right) & \text{if } a+b = -1 \text{ and } a < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (52)$$

Step 3:

$$F_3(a, b) = a \int_{\mathbb{R}} \delta_{\alpha}(x-1) \delta_{\alpha}(ax-b+1) dx$$

Similarly, we solve:

$$\frac{-1+b}{a} = 1 \Leftrightarrow a-b = -1 \quad (53)$$

and we obtain the following expression:

$$F_3(a, b) = \begin{cases} \frac{3}{4\alpha} & \text{if } a = 1 \text{ and } b = 2 \\ \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a}\right) & \text{if } a-b = -1 \text{ and } a > 1 \\ \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a\right) & \text{if } a-b = -1 \text{ and } a < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (54)$$

Step 4:

$$F_4(a, b) = a \int_{\mathbb{R}} \delta_{\alpha}(x-1) \delta_{\alpha}(ax-b-1) dx$$

We solve:

$$\frac{1+b}{a} = 1 \Leftrightarrow a-b = 1 \quad (55)$$

and we have:

$$F_4(a, b) = \begin{cases} \frac{3}{4\alpha} & \text{if } a = 1 \text{ and } b = 0 \\ \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a}\right) & \text{if } a-b = 1 \text{ and } a > 1 \\ \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a\right) & \text{if } a-b = 1 \text{ and } a < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (56)$$

Step 5: back to the whole functional F

The study of each sub-functional has put forward the conditions for which two peaks are aligned. Putting back together the results leads to focusing on the conditions when both pairs of peaks are aligned at the same time. By comparing the conditions on a and b for each sub-functional, the only simultaneous association is between F_1 and F_4 , with $a = 1$ and $b = 0$. In this case, we have $F(1, 0) = F_1(1, 0) + F_4(1, 0) = \frac{3}{4\alpha} + \frac{3}{4\alpha} = \frac{3}{2\alpha}$.

Conclusion :

$$F(a, b) = \begin{cases} \frac{3}{2\alpha} & \text{if } a = 1 \text{ and } b = 0, \\ \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a} \right) & \text{if } a > 1 \text{ and } \begin{cases} a + b = 1 \\ \text{or } a + b = -1 \\ \text{or } a - b = 1 \\ \text{or } a - b = -1 \end{cases} \\ \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a \right) & \text{if } a < 1 \text{ and } \begin{cases} a + b = 1 \\ \text{or } a + b = -1 \\ \text{or } a - b = 1 \\ \text{or } a - b = -1 \end{cases} \\ 0 & \text{otherwise.} \end{cases} \quad (57)$$

To conclude, we need to show that $F(a, b) \leq \frac{3}{2\alpha}$ so that the couple $(a, b) = (1, 0)$ is the optimum. To this aim, we focus on:

$f_1(a) = \frac{1}{2\alpha} \left(1 + \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a} \right)$ for $a > 1$ and $f_2(a) = \frac{1}{2\alpha} \left(a + \frac{1}{\pi(1+a)(1-a)} \sin \pi a \right)$ for $0 < a < 1$ and we split the proof on three parts:

- $0 < a < 1$

In order to show that $f_2(a) \leq \frac{3}{2\alpha}$, we need to assess that $\frac{1}{\pi(1+a)(1-a)} \sin \pi a \leq 2$. We have:

$$\begin{aligned} & \frac{1}{\pi(1+a)(1-a)} \sin \pi a \\ &= \frac{1}{\pi(1+a)(1-a)} \sin \pi(1-a) \\ &\leq \frac{1}{(1+a)} \leq 1 \end{aligned} \quad (58)$$

on the interval $[0; 1]$.

- $1 < a < 2$

In order to show that $f_1(a) \leq \frac{3}{2\alpha}$ on the interval $[1; 2]$, we need to verify that

$$\begin{aligned} & \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a} \leq 2. \text{ We have:} \\ & \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a} \\ &= \frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi(a-1)}{a} \\ &\leq \frac{a^3}{\pi(a+1)(a-1)} \times \frac{\pi(a-1)}{a} \\ &\leq \frac{a^2}{a+1} \leq \frac{4}{3} \end{aligned} \quad (59)$$

on the interval $[1; 2]$.

- $a > 2$

$$\frac{a^3}{\pi(a+1)(a-1)} \sin \frac{\pi}{a} \leq \frac{a^2}{a^2-1} \leq \frac{4}{3} \quad (60)$$

on the interval $[2; +\infty]$.

This concludes the proof by showing that the functional has a unique global maximizer for the sought parameters $(a, b) = (1, 0)$.

References

1. Allaire, S., Kim, J.J., Breen, S.L., Jaffray, D.A., Pekar, V.: Full orientation invariance and improved feature selectivity of 3d sift with application to medical image analysis. In: Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on, pp. 1–8. IEEE (2008)
2. Aubert, G., Kornprobst, P.: Mathematical problems in image processing: partial differential equations and the calculus of variations, vol. 147. Springer (2006)
3. Bay, H., Ess, A., Tuytelaars, T., Gool, L.V.: SURF: Speeded up robust features. Computer Vision and Image Understanding (CVIU) **110(3)**, 346–359 (2008)
4. Brown, L.: A survey of image registration techniques. ACM Computing Surveys **24(4)**, 325–376 (1992)
5. Buades, A., Coll, B., Morel, J.M.: A review of image denoising algorithms, with a new one. Multiscale Modeling and Simulation **4(2)**, 490–530 (2005)
6. Dellinger, F., Delon, J., Gousseau, Y., Michel, J., Tupin, F.: SAR-SIFT: A SIFT-like algorithm for applications on sar images. In: Geoscience and Remote Sensing Symposium (IGARSS), 2012 IEEE International, pp. 3478–3481. IEEE (2012)
7. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24(6)**, 381–395 (1981)
8. Haber, E., Modersitzki, J.: Intensity gradient based registration and fusion of multi-modal images. Medical image computing and computer-assisted intervention: MICCAI **46**, 292–299 (2006)
9. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge university press (2003)
10. Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, S.M.,

- Schnabel, J.A.: Mind: Modality independent neighbourhood descriptor for multimodal deformable registration. *Medical Image Analysis* **16**(7), 1423–1435 (2012)
11. Hermosillo, G., Chefd’Hotel, C., Faugeras, O.: Variational methods for multimodal image matching. *International Journal of Computer Vision* **50**(3), 329–343 (2002)
 12. Irani, M., Anandan, P.: Robust multi-sensor image alignment. In: *Computer Vision, 1998. Sixth International Conference on*, pp. 959–966. IEEE (1998)
 13. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2), 91–110 (2004)
 14. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging* **16**(2), 187–198 (April 1997)
 15. Mann, S., Picard, R.W.: Video orbits of the projective group a simple approach to featureless estimation of parameters. *Image Processing, IEEE Transactions on* **6**(9), 1281–1295 (1997)
 16. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**(10), 1615–1630 (2005)
 17. Odobez, J.M., Bouthemy, P.: Robust multi-resolution estimation of parametric motion models. *Journal of visual communication and image representation* **6**(4), 348–365 (1995)
 18. Paganelli, C., Peroni, M., Pennati, F., Baroni, G., Summers, P., Bellomi, M., Riboldi, M.: Scale invariant feature transform as feature tracking method in 4D imaging: A feasibility study. In: *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 6543–6546 (2012)
 19. Pluim, J., Maintz, J., Viergever, M.: Image registration by maximization of combined mutual information and gradient information. *IEEE Transactions on medical imaging* **19**(8), 809–814 (2000)
 20. Pratt, W.: *Digital Image Processing*. John Wiley and Sons, Inc., NY (1978)
 21. Roshni, V., Revathy, K.: Using mutual information and cross correlation as metrics for registration of images. *Journal of Theoretical & Applied Information Technology* **4**(6) (2008)
 22. Spontón, H., Cardelino, J.: A review of classic edge detectors. In: *Image Processing On Line* (2012)
 23. Sun, Y., Jolly, M.P., Moura, J.F.: Integrated registration of dynamic renal perfusion MR images. In: *ICIP*, pp. 1923–1926 (2004)
 24. Toews, M., Zöllei, L., Wells, W.M.: Feature-based alignment of volumetric multi-modal images. In: *Information Processing in Medical Imaging*, pp. 25–36. Springer (2013)
 25. Viola, P., Wells III, W.M.: Alignment by maximization of mutual information. *International journal of computer vision* **24**(2), 137–154 (1997)
 26. Yu, L., Zhang, D., Holden, E.J.: A fast and fully automatic registration approach based on point features for multi-source remote-sensing images. *Computers & Geosciences* **34**(7), 838–848 (2008)
 27. Zitova, B., Flusser, J.: Image registration methods: A survey. *Image and Vision Computing* **21**, 977–1000 (2003)