

Estimation of an Observation Satellite's Attitude Using Multimodal Pushbroom Cameras

Régis Perrier, Elise Arnaud, Peter Sturm, and Mathias Ortner

Abstract—Pushbroom cameras are widely used for earth observation applications. This sensor acquires 1D images over time and uses the straight motion of the satellite to sweep out a region of space and build a 2D image. The stability of the satellite is critical during the pushbroom acquisition process. Therefore its attitude is assumed to be constant over time. However, the recent manufacture of smaller and lighter satellites to reduce launching cost has weakened this assumption. Small oscillations of the satellite's attitude can result in noticeable warps in images, and geolocation information is lost as the satellite does not capture what it ought to. Current solutions use inertial sensors to control the attitude and correct the images, but they are costly and of limited precision. As the warped images do contain information about attitude variations, we suggest using image registration to estimate them. We exploit the geometry of the focal plane and the stationary nature of the disturbances to recover undistorted images. We embed the estimation in a Bayesian framework where image registration, a prior on attitude variations and a radiometric correction model are fused to retrieve the motion of the satellite. We illustrate the performance of our algorithm on four satellite datasets.

Index Terms—Multimodal image registration, satellite attitude, pushbroom cameras, hyperparameter learning, maximum a posteriori estimator

1 INTRODUCTION

THE pushbroom camera is widely used since the 30's for earth observation applications on board of a flying vehicle like a satellite or an aircraft; the article [10] provides a good historical review of this sensor. Its specific design is dedicated to exploit the motion of the vehicle; this linear camera which captures several 1D images—or lines of pixels—over time, and sweeps out a region of space as the imager moves straight and orthogonally to its acquisition axis. The result of gathering the 1D images side by side is an infinite image strip which can be cut to form 2D images. The acquisition process is summarized in Fig. 1.

This sensor is manufactured in large quantities for daily use in document scanners, fax machines or bar code readers; we will focus here on satellite imagery. It has several advantages over the conventional pinhole camera like: a very high resolution image recording which can produce gigapixel 2D images in few seconds [23], an optimal use of the motion of the satellite to capture images without redundancy [34], and a low production cost with a high robustness. Note that the electronic

equipment of a satellite needs to be hardened to endure space environment, thus a specific design of the charge-coupled device is necessary and none of the classical digital cameras available on the market could be used.

The pushbroom sensor requires a constant satellite attitude¹ during the image acquisition. This strong assumption was considered to be true several years ago [10]. However, the new generation of satellites is smaller and lighter by design to reduce launching cost and make it more easy to pilot, making oscillations of the spacecraft around its rotation axis more likely. Also, with the increase of camera resolution and lens quality, any very small dynamic disturbance of the imager, like a rotation of a few microradians, can visibly warp each 1D image, and consequently the whole 2D image. Fig. 2 is a toy example of the warps one could observe on a chessboard with an oscillating pushbroom acquisition; the picture looks wobbly in an almost similar manner as an image acquired with a rolling shutter [28].

An observation satellite usually has several pushbroom cameras set in parallel on its focal plane to record multispectral images. Fig. 3 is a schematic view of such focal plane; it shows that a point on the earth is observed by each camera at different time instants, which depend on the spacing between cameras and the speed of the satellite. Thus, in order to build a color image as a superposition of the red, blue, and green channels, it is necessary to first register those channels in a common scene reference. This is an easy step as the focal plane's configuration is known, but one may see that any attitude variations of the focal plane during the acquisition process leads to a tricky problem. In this case, all the images are warped by the same temporal process, but those warps are shifted in time according to the spatial locus of the sensors.

1. Orientation of the vehicle in a 3D space defined by the three rotations: yaw, roll, and pitch.

• R. Perrier is with the CEA Leti, Grenoble, France.

E-mail: perrier.regis@gmail.com.

• E. Arnaud is with the Université Joseph Fourier, Grenoble, France.

E-mail: elise.arnaud@inria.fr.

• P. Sturm is with the INRIA, Grenoble, France.

E-mail: peter.sturm@inria.fr.

• M. Ortner is with the EADS Astrium, Toulouse, France.

E-mail: mathias.ortner@astrium.eads.net.

Manuscript received 10 Dec. 2013; revised 14 Sept. 2014; accepted 18 Sept. 2014. Date of publication 25 Sept. 2014; date of current version 4 Apr. 2015.

Recommended for acceptance by M. S. Brown.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2014.2360394

1. Orientation of the vehicle in a 3D space defined by the three rotations: yaw, roll, and pitch.




Fig. 1. Pushbroom acquisition principle: the camera is moving straight along the y axis and recording 1D images over time denoted by t . x is the camera axis and z the orthogonal axis to the image plane. The attitude of the camera is defined by the yaw (rotation about z), the roll (rotation about y), and the pitch (rotation about x).

The first consequence of attitude variations of the satellite is the bleeding artefacts on the edges of the recomposed color image, but the major problem is the loss of geolocation information as the satellite does not capture the scene it is programmed for. Most of the research and industrial use of satellite images needs accurate data and cannot afford such defects. Thus, those attitude variations have to be estimated to correct the images.

This paper presents an original method to find the attitude variations of a satellite based on the registration of its acquired images. This is in contrast to the vast majority of satellite attitude estimation methods which use external sensors like inertial measurement units (IMUs). We first review related works to motivate the originality of this paper. Then, we describe how two images acquired by the satellite can be registered into a common spatial reference. To do so, we exploit the linear geometry and the positions of the pushbroom cameras on the focal plane. This principle can be applied to any pair of images coming from pushbroom cameras of the satellite focal plane, it is a main contribution of the paper. However, this problem is ill-posed and needs some specific prior on the attitude variations to be solved. The fourth section explains the embedding of the estimation process in a Bayesian framework where a prior




Fig. 2. Example of warps in a regular checkerboard when the pushbroom camera is tilting around its three rotation axes.




Fig. 3. Standard focal plane geometry of an observation satellite with four pushbroom cameras: panchromatic, blue, green, and red (respectively enumerated as 1, 2, 3, and 4). What is seen by camera 2 at time t will be seen by camera 1 at time $t + \delta$.

on the attitude variations, an image registration term, and a radiometric correction model between the two images are combined. It also describes two estimation algorithms based on the Lucas Kanade technique [2], [25]. To insure an accurate multimodal image registration process, we design a radiometric model at a pixel scale which, according to the experiments, achieves a matching error lower than a tenth of a pixel. This is another contribution of this paper. The next section shows how the Bayesian framework allows to automatically adjust the statistical parameters of the model, which will be called hyperparameters, with an optimization procedure of the model's evidence [27]. This makes our method almost completely free of parameter tuning. We also investigate the integration of the radiometric coefficients in the complete model, which lets us estimate the attitude variations without the need to estimate the radiometric coefficients. This, to our best knowledge, has not yet been formalized and tested for multimodal image registration purposes. Finally, we evaluate the performances of our solution and compare it over standard approaches with four satellite datasets given by EADS Astrium.²

2 RELATED WORK

Unlike the pinhole model, little attention was paid to the pushbroom case; these cameras are less known and datasets are harder to obtain. Concerning their calibration, the corner stone is probably the work of Gupta and Hartley [13], [14]. They derive two models to explain correspondences between 2D and 3D points, they prove the existence of an essential matrix for pushbroom cameras and show reconstruction results of a 3D surface with satellite dataset. More recently, a linear calibration method that uses homographies induced by images of a planar object has been presented in [8]. However, both methods assume that the attitude is fixed during the acquisition, also they are not suited for the case of several pushbroom cameras belonging to the same focal plane.

Airborne imagery with helicopters or airplanes also uses pushbroom cameras and has to deal with significant attitude variations because of the wind. Nevertheless, it can rely on very accurate IMUs on board to estimate the attitude over time and rectify the images [11]. Thus, people focussed

2. EADS Astrium is a European aerospace company and satellite manufacturer.

more on IMU sensors fusion with global positioning system [5], or ground control points which are manually picked features on the images with known ground coordinates [32], to get an estimate of the attitude. This technology has lead to impressive 3D reconstruction methods of urban areas since a few years [16].

The satellite case is more complex as such inertial sensors cannot be easily brought into space. Apart from being very costly due to their specific electronic design, they can be prone to strong errors in their measurements because of the space environment. Also their low sampling rate, usually below a hundred Hertz, is not sufficient compared to the pushbroom acquisition rate of 1D images which can range between several hundred to few thousands Hertz.

As the warps on the images do contain information about attitude variations, a method for estimating them using image registration has been suggested in [6]; this is apparently the earliest reference to such an idea, a similar solution was proposed more recently in [20]. This method is a multi-step procedure for all pairs of images which first extracts corner points on one image, seeks their corresponding points in a local neighborhood in the other image with a cross correlation procedure, infers a relative attitude which is finally deconvolved to obtain an absolute attitude variation estimate for all images; this matter of relative and absolute attitude will become clearer in the next section. However, this solution has several drawbacks: it uses small correlation windows to estimate the local translation between a pair of images whereas the warp is known to be rigid along a line of pixels due to the geometry of the pushbroom camera. It has several parameters which need to be manually tuned like the sensitivity of the corner detector, the size of the correlation window or the regularization term of the deconvolution step. It offers no easy way to add sensor measurements or priors on the attitude variations to the estimation process. Its performances depend on the observed scene; for example desert regions will have less feature points to match, whereas sea surfaces will likely produce false features matching because of the repeated wave patterns which move over time.

Image registration, as a key preprocessing step in many computer vision applications [40], is a long-standing research field; many solutions exist but none can really satisfy alone this task in all its forms. Instead, one needs to specify the nature of its image registration problem to get the best performances: whether the warps have a parametric form or not—[39] and [41] are among the best reviews for those two cases—if the misalignment is large—in which case feature based techniques should be favored [24]—or of a few pixels [25], and if the sensors capture the same radiometric modalities. Other criterions exist, those are just well adapted to our purpose; [17] gives a good overview of image registration for satellite images.

In this problem where deformations are of a few pixels, direct image registration methods [19], also known as pixel-based [39], or Lucas Kanade methods [25], should be well suited: they make an optimal use of all the pixels in the images, they have a sub-pixel accuracy which is needed in this application, and finally they are fast and robust if the assumption of small warps is not broken. However, the sum of square differences (SSDs) criterion of the Lucas

Kanade is very sensitive to outliers such as radiometric differences. Several solutions were proposed to deal with the multimodal case: taking a more robust similarity criterion like Mutual Information [7] but at the expense of an increased computational cost, using local correlation methods [17] but as we previously explained it goes against our will to consider a parametric warp related to the geometry of the pushbroom camera, high pass filtering the images before the registration process [18], rectifying one image with a radiometric linear model to match the other at a global scale [2] and at a local scale [21], [42], or modeling the radiometric intensity transform [12]. In our case where the radiometric differences are small, some of those methods are either too complex or may be less accurate than the basic SSD criterion. So as to make sure that the image registration performs at least as good as the SSD criterion in this case, a linear radiometric correction model at a pixel scale seems to be a good solution. The statistical relevance of this model has already been investigated on real images for multi-sensor super resolution in [42]; this last article clearly motivates our work. Moreover, we will show how it can fit efficiently in the Lucas Kanade algorithm with a two step optimization procedure. There is however a strong risk of overfitting with this solution [21], thus specific priors for the radiometric coefficients need to be designed as well as a method to infer the parameters of those priors. This is also a contribution of this paper over the state of the art in multimodal image registration.

3 UNDERSTANDING IMAGE REGISTRATION

Fig. 1 summarises the acquisition principle of a single pushbroom camera, now let's consider the case of two pushbroom sensors set in parallel on a focal plane. We will make the following assumptions:

- The distance between the earth and the satellite is large so that the observed scene can be considered as being planar; this implies that parallax effects between the lines of sight of the two cameras are negligible. In practice, the use of digital elevation model (DEM) as described later will allow to relax this hypothesis.
- The satellite moves straight and at constant speed at the time scale of a 2D image acquisition.
- Given the small time shift between the observation of the same scene point by the two cameras, occlusion artefacts are unlikely to happen and warps on images cannot exceed a few pixels.

Let $\mathcal{I}$ and $\mathcal{J}$ be the image functions of respectively camera i and j where i and $j \in [1, 4]$, and $j \neq i$ according to Fig. 3. Given those hypothesis, we expect to have the following relationship between the two images recorded by the pushbroom cameras³:

$$\mathcal{I}(x, t) - \mathcal{J}(x, t - \delta) \sim \mathbf{N}(0, \sigma_{\mathcal{I}}^2), \quad (1)$$

3. We use the following math notations: column vectors are bold (e.g., $\mathbf{x}$), scalars are written in italics (e.g., x), matrices are upper case letters (e.g., $\mathbf{X}$), calligraphic font denotes functions (e.g., $\mathcal{X}$), $\mathbf{N}(\mu, \sigma^2)$ is the Gaussian probability distribution with mean μ and standard deviation σ .

where $x \in [1, n]$ is the discrete pixel position along the camera axis, $t \in [1, \tau]$ is the discrete acquisition time index, and δ is the real time shift value in pixel unit according to the focal plane. Images are defined on a discrete grid of pixels, but image intensity values on a non integer pixel location can be obtained through interpolation techniques. Literally, this expression means that one expects to capture at time t with camera i what has been recorded by camera j at time $t - \delta$, up to sensor noise and radiometric sensitivity differences between both cameras, which are represented as a zero mean Gaussian distribution with standard deviation $\sigma_{\mathcal{I}}$. During the acquisition process, small attitude variations of the satellite can occur and warp the images according to a function which will be denoted by $\mathcal{W}$. Let θ_t be the attitude variation of the satellite over time; for the sake of clarity we will assume that this variable could either be the yaw, the roll or the pitch angle. We will make the following hypotheses:

- There exists a function $\mathcal{W}$ which, given a pixel location (x, t) on an image and an attitude variation θ_t , outputs a new pixel location on the image:

$$\mathcal{W} : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2 \quad (2)$$

$$\mathcal{W}(x, t; \theta_t) = (x', t'). \quad (3)$$

- This function is invertible and possesses the following composition rule:

$$\mathcal{W}(\mathcal{W}(x, t; \theta_t); \theta'_t) = \mathcal{W}(x, t; \theta_t + \theta'_t). \quad (4)$$

- When $\theta_t = 0$, the function is the identity:

$$\mathcal{W}(x, t; 0) = (x, t). \quad (5)$$

Regarding this function, it is hard to obtain an analytical form for $\mathcal{W}$ as it depends on the cinematic of the satellite as well as the ground elevation of the observed scene. If it is relatively easy to project the line of sight of the pushbroom camera on the ground given a DEM, the inverse operation which computes the location of a point on the pushbroom image given its spatial coordinates on the ground is insoluble without strong hypotheses like in [31]. Function $\mathcal{W}$ should be seen as an operator which explains how pixels move in an image when the satellite is not steady during the acquisition process. To the first order, the roll and the pitch are respectively translations of the 1D image along the x and t axis, as can be seen in Fig. 2. The experiments section will give more insight into this function, but any theoretical development of $\mathcal{W}$ is beyond the scope of this paper. In a non steady case of the satellite, the left part of Equation (3) becomes:

$$\mathcal{I}(\mathcal{W}(x, t; \theta_t)) - \mathcal{J}(\mathcal{W}(x, t - \delta; \theta_{t-\delta})), \quad (6)$$

where one should note the different time indices t and $t - \delta$ between the attitude variations of $\mathcal{I}$ and $\mathcal{J}$ which are linked to the focal plane geometry. Following our hypotheses and relation (4), we have:

$$\mathcal{I}(x, t) - \mathcal{J}(\mathcal{W}(x, t - \delta; \theta_t^\delta)), \quad (7)$$




Fig. 4. Power spectrum of k^δ for several values of δ .

where:

$$\theta_t^\delta = \theta_t - \theta_{t-\delta}. \quad (8)$$

One can look for θ_t^δ with the following expression:

$$\hat{\theta}_t^\delta = \underset{\theta_t^\delta}{\operatorname{argmin}} \sum_x (\mathcal{I}(x, t) - \mathcal{J}(\mathcal{W}(x, t - \delta; \theta_t^\delta)))^2. \quad (9)$$

This criterion is optimal with respect to the statistics of Equation (1), it is the maximum likelihood estimator of θ_t^δ and the maximized function is also well known as the SSD energy [39]. One can understand that image registration between $\mathcal{I}$ and $\mathcal{J}$ provides a relative estimate of the attitude variation θ_t^δ as opposed to an absolute one with θ_t directly. Indeed, trying to estimate θ_t from Equation (9) without the knowledge of $\theta_{t-\delta}$ is an ill-posed problem. One should first estimate $\theta_{t-\delta}$, but one would need to know $\theta_{t-2\delta}$ prior to this step, and this becomes an endless loop. So as to get a better understanding of the ill-posed nature of the system, let us here observe that it can be reformulated as a deconvolution problem where:

$$\theta_t^\delta = (k^\delta \star \theta)_t, \quad (10)$$

with $\star$ being the convolution operator and k^δ being a discrete convolution kernel⁴ which differentiates over time the attitude variations θ_t according to the value of δ . It is now interesting to look at the spectrum of k^δ in Fig. 4 with respect to the value of δ , it tells where the frequency content of θ_t is lost after the convolution process. Quite simply from the shape of the spectrums, each place where the black curve drops down is a loss of information about θ_t . Yet according to the value of δ , which we remind as being dependent on the focal plane geometry, the deconvolution problem will be more or less ill-posed with respect to the spectrum of θ_t . It should be clear that a prior information on θ_t is needed to estimate the absolute attitude variations from the images. Another important issue is the different radiometric sensitivity of the cameras. As already noted, Equation (1) should be understood as up to radiometric effects between both cameras. If the spectral channels generally overlap between all the pushbroom cameras on a satellite focal plane, which makes the images look very similar,

4. Let $\delta = 3$, the associated discrete kernel will be $k^\delta = [-1 \ 0 \ 0 \ 1]^T$. Note that any non integer value of δ can also be linearly approximated with a kernel, for example $\delta = 2.5$ and $k^\delta = [-1 \ 0 \ 0.5 \ 0.5]^T$.

it is also well known that the SSD criterion of Equation (9) is not robust to any small bias or outliers in the images. Several solutions were described in the previous section to deal with multimodal image registration. As our major concern is the accuracy of the attitude variation estimate, we will use a model at a pixel scale, and restate the problem as:

$$a_{xt} + b_{xt} \mathcal{I}(x, t) - \mathcal{J}(\mathcal{W}(x, t - \delta; \theta_t^\delta) \sim \mathbf{N}(0, \sigma_{\mathcal{I}}^2), \quad (11)$$

where a_{xt} and b_{xt} are radiometric coefficients which rectify any offset and spectral sensitivity differences between the two images. In this case, the unknowns are a_{xt} , b_{xt} , and θ_t for all x and t ; it is an expensive solution as it increases the number of unknowns at a rate of twice the number of pixels used in an image for the registration process. Moreover the problem is still ill-posed as one of the trivial solutions is to set b_{xt} and θ_t to 0, and let $a_{xt} = \mathcal{J}(x, t)$, therefore we need again to define priors over the radiometric coefficients.

4 MAXIMUM A POSTERIORI (MAP) FORMULATION

An easy way to deal with sensor measurements and priors over the unknowns may be to see this problem in a Bayesian perspective. Let θ , $\mathbf{a}$, $\mathbf{b}$, $\mathbf{i}$, and $\mathbf{j}$ be respectively the vectorized forms of the attitude variations, the radiometric coefficients, and the images $\mathcal{I}$ and $\mathcal{J}$. What we seek to maximize is the joint probability of all the variables, which writes as follows:

$$\begin{aligned} p(\theta, \mathbf{a}, \mathbf{b}, \mathbf{i}, \mathbf{j}) &= p(\theta, \mathbf{a}, \mathbf{b} | \mathbf{i}, \mathbf{j}) p(\mathbf{i}, \mathbf{j}) \\ &= \frac{p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) p(\theta, \mathbf{a}, \mathbf{b})}{p(\mathbf{i}, \mathbf{j})} p(\mathbf{i}, \mathbf{j}) \quad (12) \\ &= p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) p(\theta) p(\mathbf{a}) p(\mathbf{b}). \end{aligned}$$

The first line is a simple multiplication rule of probabilities, the second one uses Bayes theorem, and the third one assumes independence between the radiometric coefficients and the attitude variations such that $p(\theta, \mathbf{a}, \mathbf{b}) = p(\theta) p(\mathbf{a}) p(\mathbf{b})$. Thus, the optimization problem we want to solve is:

$$\hat{\theta}, \hat{\mathbf{a}}, \hat{\mathbf{b}} = \underset{\theta, \mathbf{a}, \mathbf{b}}{\operatorname{argmax}} p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) p(\theta) p(\mathbf{a}) p(\mathbf{b}), \quad (13)$$

which is also known as the maximum a posteriori estimator. This expression brings out four terms: the likelihood probability of the images $p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b})$ which explains the link between the measurements and the variables of interest, and three prior probabilities about which we will talk later.

4.1 Likelihood of Images

In probabilistic terms it could be understood as a measure of "how likely are the measurements $\mathbf{i}$ and $\mathbf{j}$ if we were given θ , $\mathbf{a}$ and $\mathbf{b}$?" The previous section has indirectly talked about the likelihood term in Equation (11), expressing its probabilistic yields:

$$p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) = \prod_{t=\delta+1}^{\tau} \prod_{x=1}^n p(\mathcal{I}(x, t), \mathcal{J}(x, t) | \theta_t, a_{xt}, b_{xt}), \quad (14)$$

where it is assumed that the noise in the images is independently distributed over the pixels space. Let M_i be a matrix of size $(n_p \times n_p)$ with the image vector $\mathbf{i}$ on its diagonal and

zero values elsewhere, and $n_p = n(\tau - \delta)$ the total number of pixels considered in the registration process. Given Equation (11), the likelihood probability writes as:

$$p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) = c_{\mathcal{I}} \exp(-0.5 \sigma_{\mathcal{I}}^{-2} \|\mathbf{a} + M_i \mathbf{b} - \mathbf{j}_{\theta}\|^2), \quad (15)$$

with:

$$c_{\mathcal{I}} = (2\pi \sigma_{\mathcal{I}}^2)^{-\frac{n(\tau-\delta)}{2}}. \quad (16)$$

Image $\mathbf{j}$ is indexed with θ to remind its dependence on the attitude variations, like matrix M_i with image $\mathbf{i}$; $\|\cdot\|$ is the vector norm such that $\|\mathbf{x}\| = \sqrt{\mathbf{x}^T \mathbf{x}}$.

4.2 Attitude Variations Prior

In order to design a specific prior for the variable θ_t , one needs to wonder about the origin of those oscillations of the satellite around its rotation axis. Attitude variations of the satellite can originate from its engines which are vibrating, but also from space turbulences which make it deviate from its trajectory. Also, the physical state of the satellite is of great importance as its angular momentum can be quite different when its solar panels are spread out or not, this could impact amplitude and frequency contents of the oscillations. Thus, choosing an informative prior which could work on a long time period is difficult and may bias the estimation process if it is too specific. However, due to the strong inertia of the satellite, sudden motion of the spacecraft is unlikely to happen such that smoothness over time is a trustworthy property of the attitude variations. Therefore, we will use a one order Markov model to explain how attitude variations change over time:

$$\theta_t = \theta_{t-1} + \epsilon \quad \epsilon \sim \mathbf{N}(0, \sigma_{\epsilon}^2). \quad (17)$$

This is like an autoregressive model (AR) with just one coefficient equal to 1. It is possible to design higher order models [29], [30], but learning the autoregressive coefficients is challenging. Moreover, it is difficult to understand how well such a model generalizes during the lifetime of a satellite.

One can express the attitude prior with Equation (17):

$$p(\theta) = p(\theta_1) \prod_{t=2}^{\tau} p(\theta_t | \theta_{t-1}) \quad \text{where} \quad \theta_1 \sim \mathbf{N}(0, \sigma_{\theta_1}^2). \quad (18)$$

We call D the $(\tau \times \tau)$ matrix which differentiates in time the attitude vector θ and let Σ_{θ} be the $(\tau \times \tau)$ diagonal matrix with $\sigma_{\theta_1}^2$ being the first element of the diagonal, and σ_{θ}^2 elsewhere. Equation (18) gives in vector form:

$$p(\theta) = c_{\theta} \exp(-0.5 \|D\theta\|_{\Sigma_{\theta}}^2), \quad (19)$$

with:

$$c_{\theta} = (2\pi \sigma_{\theta_1}^2)^{-0.5} (2\pi \sigma_{\theta}^2)^{-\frac{\tau-1}{2}}. \quad (20)$$

The operator $\|\cdot\|_{\Sigma}$ is the weighted vector norm such that $\|\mathbf{x}\|_{\Sigma} = \sqrt{\mathbf{x}^T \Sigma^{-1} \mathbf{x}}$.

There is no way to retrieve the offset information for θ_t because of the convolution nature of the problem. The role of $p(\theta_1)$ is to keep the first value of the attitude around zero,

and avoid any ill conditioning of the system of equations to solve; the parameter σ_{θ_1} should be set to a sufficiently large value to give some freedom to the optimization process.

4.3 Radiometric Coefficients Prior

We do not have much insight on the nature of the radiometric coefficients $\mathbf{a}$ and $\mathbf{b}$, except that they should bring a minor correction to the image registration process, and in order to relate Equations (7) and (11), coefficients a_{xt} and b_{xt} should be respectively close to 0 and 1. An interesting point raised by the image processing community is that the vast majority of image registration techniques needs high frequency features in images to make correspondences like corners [15], lines [4], image derivatives [18], or higher level features [24]. In our case where images are very similar and occlusion is unlikely to happen, we could let the radiometric differences be low frequency components in the images. The idea behind is to make the high frequency features of the images drive the registration process while letting the disturbing radiometric features be smoothed over the images.

Considering spatial smoothness on both dimensions x and t for the coefficients $\mathbf{a}$:

$$\begin{aligned} a_{xt} - a_{xt-1} &\sim p(a_{xt} | a_{xt-1}) = \mathbf{N}(0, \sigma_{a_s}^2) \\ a_{xt} - a_{x-1t} &\sim p(a_{xt} | a_{x-1t}) = \mathbf{N}(0, \sigma_{a_s}^2), \end{aligned} \quad (21)$$

we can define the following prior:

$$\begin{aligned} p(\mathbf{a}) &= p(a_{n\tau}) \prod_{t=\delta+2}^{\tau} \prod_{x=1}^n p(a_{xt} | a_{xt-1}) \\ &\prod_{t=\delta+1}^{\tau} \prod_{x=2}^n p(a_{xt} | a_{x-1t}) \quad \text{where } a_{n\tau} \sim \mathbf{N}(0, \sigma_{a_{n\tau}}^2). \end{aligned} \quad (22)$$

We denote by R the matrix which differentiates along x and t the vector of coefficients $\mathbf{a}$, and by Σ_a the diagonal matrix⁵ which contains $\sigma_{a_{n\tau}}^2$ as its last element and $\sigma_{a_s}^2$ elsewhere. This gives in vector form:

$$p(\mathbf{a}) = c_a \exp(-0.5 \|\mathbf{Ra}\|_{\Sigma_a}^2), \quad (23)$$

with:

$$c_a = (2\pi\sigma_{a_{n\tau}}^2)^{-\frac{1}{2}} (2\pi\sigma_{a_s}^2)^{-\frac{n(2\tau-2\delta-1)-\tau+\delta}{2}}. \quad (24)$$

The parameter σ_{a_s} tunes the amplitude of a coefficient with respect to its neighbors along x and t , whereas $\sigma_{a_{n\tau}}$ sets the amplitude of coefficient $a_{n\tau}$ alone. This last coefficient still has a global impact on the model as all the coefficients are connected together.

The prior for $\mathbf{b}$ will be the same, but with a different smoothness parameter σ_{b_s} , and a different prior probability distribution for $b_{n\tau}$ which will be $\mathbf{N}(1, \sigma_{b_{n\tau}}^2)$, according to the observation we made at the beginning of the paragraph. Hence we have:

$$p(\mathbf{b}) = c_b \exp(-0.5 \|\mathbf{Rb} - \boldsymbol{\mu}_b\|_{\Sigma_b}^2), \quad (25)$$

with $\boldsymbol{\mu}_b$ being the mean vector which elements equal zero, except for its last term which corresponds to the mean of $p(b_{n\tau})$ and equals 1, and:

$$c_b = (2\pi\sigma_{b_{n\tau}}^2)^{-\frac{1}{2}} (2\pi\sigma_{b_s}^2)^{-\frac{n(2\tau-2\delta-1)-\tau+\delta}{2}}. \quad (26)$$

4.4 MAP Estimator

Given the likelihood of images in Equation (15) and the priors in Equations (19), (23), and (25), we can express the full MAP estimator of Equation (13) in logarithmic form and up to constant values which do not depend on the variables of interest:

$$\begin{aligned} \hat{\boldsymbol{\theta}}, \hat{\mathbf{a}}, \hat{\mathbf{b}} &= \underset{\boldsymbol{\theta}, \mathbf{a}, \mathbf{b}}{\operatorname{argmax}} \quad -\sigma_{\mathcal{I}}^{-2} \|\mathbf{a} + M_1 \mathbf{b} - \mathbf{j}_{\theta}\|^2 \\ &\quad - \|\mathbf{D}\boldsymbol{\theta}\|_{\Sigma_{\theta}}^2 - \|\mathbf{Ra}\|_{\Sigma_a}^2 - \|\mathbf{Rb} - \boldsymbol{\mu}_b\|_{\Sigma_b}^2, \end{aligned} \quad (27)$$

where $\sigma_{\mathcal{I}}$, σ_{a_s} , $\sigma_{a_{n\tau}}$, σ_{b_s} , $\sigma_{b_{n\tau}}$, and σ_{θ} are parameters of the model often known as hyperparameters [22], [26], [33]. As a sum of squared residuals, this system can be efficiently solved for example with a Gauss Newton (GN) algorithm. There are several motivations for this choice. The warps between the two images and the amplitude of the radiometric coefficients are small enough such that setting $\boldsymbol{\theta} = \mathbf{0}$, $\mathbf{a} = \mathbf{0}$ and $\mathbf{b} = \mathbf{1}$ is already close to the good solution; in this case the GN algorithm leads to a fast convergence. It does not need to compute any hessian matrix of the system as it makes use of its jacobian matrix to approximate it, and moreover the jacobian matrix can be easily computed for all the prior terms. Finally, it has been shown that the GN search strategy is among the best methods for pixel based image registration in the case of small warps [2]. It was originally proposed in [25] for the task of image registration and is now well known as the Lucas Kanade algorithm. The main motivation is to approximate the warped image, $\mathbf{j}_{\theta}$ in our case, with its first order Taylor expansion:

$$\mathbf{j}_{\theta} \simeq \mathbf{j} + \underbrace{\begin{bmatrix} \frac{\partial \mathbf{j}_{\theta}}{\partial x} & \frac{\partial \mathbf{j}_{\theta}}{\partial t} \end{bmatrix}}_{\nabla_{\mathbf{x}t}^{\mathbf{j}_{\theta}}} \underbrace{\begin{bmatrix} \frac{\partial \mathcal{W}_x}{\partial \theta} \\ \frac{\partial \mathcal{W}_t}{\partial \theta} \end{bmatrix}}_{\nabla_{\theta}^{\mathcal{W}}} \boldsymbol{\theta}^{\delta}, \quad (28)$$

where $[\cdot]$ is a matrix concatenation operator, $\nabla_{\mathbf{x}t}^{\mathbf{j}_{\theta}}$ is a $(n_p \times 2n_p)$ matrix which contains partial derivatives⁶ of the warped image $\mathcal{J}$ with respect to x and t , and $\nabla_{\theta}^{\mathcal{W}}$ is the $(2n_p) \times (\tau - \delta)$ matrix of partial derivatives of the warp function $\mathcal{W}$ with respect to θ . One should remind that θ is convolved by the kernel k^{δ} as explained in Section 3. This linear operation results in the vector $\boldsymbol{\theta}^{\delta}$ whose length is $\tau - \delta$. The algorithm can be summarized by the following steps:

global algorithm

- Initialize variables $\boldsymbol{\theta} = \mathbf{0}$, $\mathbf{a} = \mathbf{0}$ and $\mathbf{b} = \mathbf{1}$ and set hyperparameters $(\sigma_{\mathcal{I}}, \sigma_{a_s}, \sigma_{a_{n\tau}}, \sigma_{b_s}, \sigma_{b_{n\tau}}, \sigma_{\theta})$.
- Do:

5. The size of matrix R is $(n(2\tau - 2\delta - 1) - \tau + \delta + 1) \times (n_p)$ and accordingly the size of matrix Σ_a is $(n(2\tau - 2\delta - 1) - \tau + \delta + 1) \times (n(2\tau - 2\delta - 1) - \tau + \delta + 1)$.

6. Those derivatives can be efficiently approximated by convolving the image with high pass filters like $[1 - 1]$ and $[1 - 1]^T$ to get, respectively, the derivatives along the x and t directions.

- 1) warp image $\mathcal{J}$ according to θ^δ to get $\mathbf{j}_\theta$,
- 2) compute $\nabla_{\mathbf{x}t}^{j_\theta}$ and $\nabla_\theta^{\mathcal{W}}$,
- 3) To get $\Delta\theta$, $\Delta\mathbf{a}$ and $\Delta\mathbf{b}$, solve:

$$\begin{aligned} \underset{\theta, \mathbf{a}, \mathbf{b}}{\operatorname{argmax}} \quad & -\sigma_{\mathcal{I}}^{-2} \|\mathbf{a} + M_1 \mathbf{b} - \mathbf{j}_\theta - \nabla_{\mathbf{x}t}^{j_\theta} \nabla_\theta^{\mathcal{W}} \theta^\delta\|^2 \\ & - \|D\theta\|_{\Sigma_\theta}^2 - \|R\mathbf{a}\|_{\Sigma_a}^2 - \|R\mathbf{b} - \boldsymbol{\mu}_b\|_{\Sigma_b}^2, \end{aligned}$$

- 4) update variables:

$$\theta = \theta + \Delta\theta, \mathbf{a} = \mathbf{a} + \Delta\mathbf{a} \quad \text{and} \quad \mathbf{b} = \mathbf{b} + \Delta\mathbf{b}$$

- while $\frac{\sqrt{\|\Delta\theta\|^2 + \|\Delta\mathbf{a}\|^2 + \|\Delta\mathbf{b}\|^2}}{n_p^2 \tau} > \xi$.

The linear expression in step 3 can be solved using the normal equations as in a linear least squares problem. The stopping criterion stops the algorithm whenever the update to the variables is too small; ξ was empirically set to 10^{-5} .

This algorithm, denominated "global" thereafter, jointly estimates attitude variations with radiometric coefficients. Yet, variables $\mathbf{a}$ and $\mathbf{b}$ could be solved analytically if θ was known as Equation (27) is linear with respect to them. The following algorithm takes advantage of this observation to split the estimation process into two steps:

two steps algorithm

- Initialize variables $\theta = 0$, $\mathbf{a} = 0$, $\mathbf{b} = 1$ and $\mathcal{J}$ to $\mathbf{j}_\theta$, set hyperparameters $(\sigma_{\mathcal{I}}, \sigma_{a_s}, \sigma_{a_{nt}}, \sigma_{b_s}, \sigma_{b_{nt}}, \sigma_\theta)$.
- Do:
 - 1) compute $\nabla_{\mathbf{x}t}^{j_\theta}$ and $\nabla_\theta^{\mathcal{W}}$,
 - 2) to get $\Delta\theta$, solve:

$$\begin{aligned} \underset{\theta}{\operatorname{argmax}} \quad & -\sigma_{\mathcal{I}}^{-2} \|\mathbf{a} + M_1 \mathbf{b} - \mathbf{j}_\theta - \nabla_{\mathbf{x}t}^{j_\theta} \nabla_\theta^{\mathcal{W}} \theta^\delta\|^2 \\ & - \|D\theta\|_{\Sigma_\theta}^2, \end{aligned}$$

- 3) update variable: $\theta = \theta + \Delta\theta$,
- 4) warp image $\mathcal{J}$ according to θ^δ to compute $\mathbf{j}_\theta$,
- 5) to get $\mathbf{a}$ and $\mathbf{b}$, solve:

$$\begin{aligned} \underset{\mathbf{a}, \mathbf{b}}{\operatorname{argmax}} \quad & -\sigma_{\mathcal{I}}^{-2} \|\mathbf{a} + M_1 \mathbf{b} - \mathbf{j}_\theta\|^2 - \|R\mathbf{a}\|_{\Sigma_a}^2 \\ & - \|R\mathbf{b} - \boldsymbol{\mu}_b\|_{\Sigma_b}^2, \end{aligned}$$

- while $\frac{\sqrt{\|\Delta\theta\|^2}}{\tau} > \xi$.

The system of equations in step 3 of the global algorithm is now approximated with two steps: step 2 estimates the attitude variations and step 5 solves a linear system to find radiometric coefficients. All variables are refined at each iteration, but whereas the first solution looks for a vector of length $\tau + 2n_p$, the two step algorithm estimates two vectors of length τ and $2n_p$. The benefits and drawbacks of this approximation will automatically come out in the experimental part where we consider the case of multiple images registration.

5 STEPPING INTO BAYES FRAMEWORK

Two questions arise from the model explained previously: how to choose the hyperparameters $\sigma_{\mathcal{I}}$, σ_{a_s} , $\sigma_{a_{nt}}$, σ_{b_s} , $\sigma_{b_{nt}}$, and σ_θ which shape the statistics of the model? As our primary focus are the attitude variations of the satellite, is there a way to drop the radiometric coefficients from the estimation process?

5.1 Hyperparameter Learning

Setting the hyperparameters is a difficult task without proper expertise on the problem or any groundtruth data which can evaluate the accuracy of the model for a given set of hyperparameters. They can be set by hand, or optimized with respect to several criteria: model complexity [1], generalization error [35], model evidence [22], [26], or photometry [3] among other methods. The number of hyperparameters to look for discards any greedy algorithm like the cross validation, and optimizing the model evidence in this case is a reliable method. The marginal likelihood probability of the images is a measure of the evidence of the model such that finding its mode with respect to the hyperparameters will hopefully select the right model. This probability distribution is as follows:

$$\begin{aligned} p(\mathbf{i}, \mathbf{j}) &= \int_{\theta, \mathbf{a}, \mathbf{b}} p(\mathbf{i}, \mathbf{j}, \theta, \mathbf{a}, \mathbf{b}) d\theta d\mathbf{a} d\mathbf{b} \\ &= \int_{\theta, \mathbf{a}, \mathbf{b}} p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) p(\theta) p(\mathbf{a}) p(\mathbf{b}) d\theta d\mathbf{a} d\mathbf{b}. \end{aligned} \quad (29)$$

Let $g(\mathbf{a}, \mathbf{b}, \theta)$ be the following function:

$$\begin{aligned} g(\mathbf{a}, \mathbf{b}, \theta) &= 0.5 \left(\sigma_{\mathcal{I}}^{-2} \|\mathbf{a} + M_1 \mathbf{b} - \mathbf{j}_\theta\|^2 \right. \\ &\quad \left. + \|D\theta\|_{\Sigma_\theta}^2 + \|R\mathbf{a}\|_{\Sigma_a}^2 + \|R\mathbf{b} - \boldsymbol{\mu}_b\|_{\Sigma_b}^2 \right), \end{aligned} \quad (30)$$

Equation (29) becomes:

$$p(\mathbf{i}, \mathbf{j}) = c_{\mathcal{I}} c_\theta c_a c_b \int_{\theta, \mathbf{a}, \mathbf{b}} \exp(-g(\mathbf{a}, \mathbf{b}, \theta)) d\theta d\mathbf{a} d\mathbf{b}, \quad (31)$$

which is analytically intractable because of the non linear nature of the g function; one needs to come to an approximation scheme to solve this problem. The Laplace approximation is a solution in this case as long as we can have a first good estimate of the attitude variations and the radiometric coefficients. Let $\tilde{\theta}$, $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ be this first good estimate, and G the hessian matrix of the function g evaluated at $\tilde{\theta}$, $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$. Equation (31) now simplifies as:

$$\begin{aligned} p(\mathbf{i}, \mathbf{j}) &\simeq c_{\mathcal{I}} c_\theta c_a c_b \exp(-g(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}, \tilde{\theta})) \\ &\quad \int_{\theta, \mathbf{a}, \mathbf{b}} \exp \left(-0.5 \begin{bmatrix} \theta - \tilde{\theta} \\ \mathbf{a} - \tilde{\mathbf{a}} \\ \mathbf{b} - \tilde{\mathbf{b}} \end{bmatrix}^T G \begin{bmatrix} \theta - \tilde{\theta} \\ \mathbf{a} - \tilde{\mathbf{a}} \\ \mathbf{b} - \tilde{\mathbf{b}} \end{bmatrix} \right) d\theta d\mathbf{a} d\mathbf{b} \\ &= c_{\mathcal{I}} c_\theta c_a c_b \exp(-g(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}, \tilde{\theta})) |2\pi G|^{-0.5}. \end{aligned} \quad (32)$$

Thus, one has to solve the following optimization problem, expressed in a logarithmic form, to find the set of hyperparameters:

$$\begin{aligned} \hat{\sigma}_{\mathcal{I}}, \hat{\sigma}_{a_s}, \hat{\sigma}_{a_{nt}}, \hat{\sigma}_{b_s}, \hat{\sigma}_{b_{nt}}, \hat{\sigma}_{\theta} = & \operatorname{argmax}_{\sigma_{\mathcal{I}}, \sigma_{a_s}, \sigma_{a_{nt}}, \sigma_{b_s}, \sigma_{b_{nt}}, \sigma_{\theta}} \log(c_{\mathcal{I}} c_{\theta} c_a c_b) \\ & - g(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}, \tilde{\theta}) - \frac{1}{2} \log|G|. \end{aligned} \quad (33)$$

All the three terms depend on the hyperparameters according to Equations (16), (20), (24), (26), and (30). In practice, the gradient of the previous equation can be computed so that a gradient based optimization method can give a solution. One needs first to estimate the attitude variations $\tilde{\theta}$ and the radiometric coefficients $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ with the algorithm as given in Section 4.4 and a set of hyperparameters initialized by hand. Then, the optimization as given in Equation (33) can be run with the first estimate and the manually tuned hyperparameters as initial values. Considering the size of the G matrix which is $(2n_p + \tau) \times (2n_p + \tau)$, hyperparameter learning has to be performed on small patches of images to be efficient.

5.2 Marginalization of Radiometric Coefficients

A major drawback of the MAP estimator as described in Section 4.4 is that one needs to optimize the attitude variations with the radiometric coefficients together to get a fine estimate, but those coefficients are not of real interest. Again, a Bayesian framework, which defines probability distributions over the variables, gives the opportunity to integrate out some of the variables we do not care about. More specifically we have:

$$\begin{aligned} p(\mathbf{i}, \mathbf{j}, \theta) &= \int_{\mathbf{a}, \mathbf{b}} p(\mathbf{i}, \mathbf{j}, \theta, \mathbf{a}, \mathbf{b}) d\mathbf{a} d\mathbf{b} \\ &= \int_{\mathbf{a}, \mathbf{b}} p(\mathbf{i}, \mathbf{j} | \theta, \mathbf{a}, \mathbf{b}) p(\theta) p(\mathbf{a}) p(\mathbf{b}) d\mathbf{a} d\mathbf{b}. \end{aligned} \quad (34)$$

One should also note that:

$$p(\mathbf{i}, \mathbf{j}, \theta) = p(\mathbf{i}, \mathbf{j} | \theta) p(\theta), \quad (35)$$

so that looking for the mode of $p(\mathbf{i}, \mathbf{j}, \theta)$ is like the MAP estimator of the attitude variations, where the radiometric model has been embedded into the equation but without the need to estimate the actual radiometric coefficients. This integral is analytically tractable because of the Gaussian nature of the linear priors for coefficients $\mathbf{a}$ and $\mathbf{b}$. As the goal is to seek the mode of the probability distribution $p(\mathbf{i}, \mathbf{j}, \theta)$, we can put it into logarithmic form and drop the terms which are constant with respect to θ . This results in the following expression to be optimized:

$$\begin{aligned} \hat{\theta} &= \operatorname{argmax}_{\theta} p(\mathbf{i}, \mathbf{j}, \theta) \\ &= \operatorname{argmax}_{\theta} -\|D\theta\|_{\Sigma_{\theta}}^2 + \mathbf{c}^T A^{-1} \mathbf{c} - \sigma_{\mathcal{I}}^{-2} \mathbf{j}_{\theta}^T \mathbf{j}_{\theta}, \end{aligned} \quad (36)$$

where:

$$A = \begin{bmatrix} R^T \Sigma_a^{-1} R + \sigma_{\mathcal{I}}^{-2} I_{\tau-\delta+n} & M_{\mathbf{i}} \\ M_{\mathbf{i}}^T & R^T \Sigma_b^{-1} R + \sigma_{\mathcal{I}}^{-2} M_{\mathbf{i}}^T M_{\mathbf{i}} \end{bmatrix} \quad (37)$$

is a sparse matrix of size $(2n_p \times 2n_p)$ and:

$$\mathbf{c} = \sigma_{\mathcal{I}}^{-2} \begin{bmatrix} I_{n_p} \mathbf{j}_{\theta} \\ M_{\mathbf{i}} \mathbf{j}_{\theta} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{n_p} \\ R^T \Sigma_b^{-1} \boldsymbol{\mu}_b \end{bmatrix}, \quad (38)$$

is a vector of length $2n_p$. Variable I_{n_p} is the identity matrix of size $(n_p \times n_p)$ and $\mathbf{0}_{n_p}$ is the null vector of length n_p . In practice, expression (36) can be solved with a gradient based optimizer; we used a Newton algorithm which needs both the gradient and the hessian of expression (36) to be computed. The term $\mathbf{j}_{\theta}^T \mathbf{j}_{\theta}$ is linearized using Equation (28), and the image $\mathcal{I}$ is constantly warped at each step of the Newton algorithm with the current attitude variations estimate, just like in the GN algorithm. Yet the matrix A must be inverted; given its large size, this is a major drawback of this method, but it does not depend on θ so the inversion only needs to be done once.

6 EXPERIMENTS

6.1 Methods

We compare several image registration methods to estimate the satellite's attitude variations. Hyperparameters learning, as described in Section 6.3, is done before their use. Those methods are as follows:

M1 Method 1 does not consider any radiometric correction model; it uses the image likelihood term as given in Equation (15) with $(\mathbf{a} = \mathbf{0}, \mathbf{b} = \mathbf{1})$ and the prior on the attitude variations as given in Equation (19) to solve:

$$\hat{\theta} = \operatorname{argmax}_{\theta} -\sigma_{\mathcal{I}}^{-2} \|\mathbf{i} - \mathbf{j}_{\theta}\|^2 - \|D\theta\|_{\Sigma_{\theta}}^2. \quad (39)$$

This equation can be efficiently solved with the GN algorithm, in the same manner as the global algorithm described in Section 4.4 but without any concern on the radiometric coefficients.

M2 Method 2 is a straightforward extension of method 1, where images are filtered with a Differences of Gaussians operator. This is done once at the beginning of the algorithm for image $\mathcal{I}$, and at each iteration for image $\mathcal{J}$. High pass filtering is a well established solution to register images of different modalities [4], [18].

M3 Method 3 estimates both the attitude variations and the radiometric coefficients, it solves Equation (27) using the global algorithm in Section 4.4. It is the complete registration model which is promoted in this paper.

M4 It solves the same equation as *M3*, but with the two step algorithm. We expect to see similar results with a reduced computational time.

M5 Method 5 marginalizes the radiometric coefficients to estimate only the attitude variations. It is described in Section 5.2 and solves Equation (36). All the previous methods use the GN search strategy; in this case, we use a Newton algorithm. To speed up its convergence, we first run *M1* to get an attitude variations estimate which is then used as an initial value for the algorithm.

M6 Method 6 has been described in the CVPR paper [30]; it uses the same likelihood term as *M1* for image registration. It does not take care of radiometric differences between images, but it has a more specific prior for the attitude variations with an autoregressive model. Hyperparameters are selected with a cross validation algorithm.

Although we have described the image registration model for a single angle of the attitude for clarity, it is better in practice to jointly estimate all angles. If θ is the 3τ attitude

vector which stacks the three angles, only few modifications need to be performed: the matrix D in Equation (19) should be replaced by a block diagonal matrix with this same matrix D replicated three times, and the matrix ∇_{θ}^W of partial derivatives of the warp function in Equation (28) should be extended along its columns to contain the three angles such that its new dimension will be $(2n_p) \times (3(\tau - \delta))$. Radiometric coefficients are linked to the images and do not depend on the number of angles we considered in the estimation process. To deal with multiple images registration in $M3$ and $M4$, one needs to solve for $n \in [1, 4]$ and $m \in [1, 4]$:

$$\hat{\theta}, \hat{\mathbf{a}}, \hat{\mathbf{b}} = \underset{\theta, \mathbf{a}, \mathbf{b}}{\operatorname{argmax}} p(\theta) \prod_{n, m; n \neq m} p(\mathbf{i}_n, \mathbf{j}_m | \theta, \mathbf{a}_{nm}, \mathbf{b}_{nm}) p(\mathbf{a}_{nm}) \cdot p(\mathbf{b}_{nm}). \quad (40)$$

This extends the two images case of Equation (13) with vectors $\mathbf{a}$ and $\mathbf{b}$ which now stack all vectors $\mathbf{a}_{nm}$ and $\mathbf{b}_{nm}$ of length n_p for each considered pair of images. Thus the number of variables to estimate linearly increases at a rate of $2n_p$ when one chooses to register two images of the focal plane in Fig. 3. In practice, we empirically chose a subset of pair of images such that $n = 1$ and $m \in [2, 4]$ to reduce the dimensionality of the problem. Literally and according to Fig. 3, the panchromatic channel (1) is set as the reference image and compared to the blue (2), green (3), and red (4) template images.

Whereas $M3$ will need to estimate a $(2\tau + 6n_p)$ vector in a single iteration of the global algorithm, $M4$ only has to infer independently a 2τ vector for the attitude (roll and pitch), and three radiometric coefficients vectors of length $2n_p$ per iteration, which greatly reduce the size of linear systems to solve.

Note that for the remaining methods, the only variable to estimate is θ but the computational load is still increased by the interpolation operation of the template images at each step.

6.2 Datasets

Each dataset has four multispectral images (panchromatic, red, blue, and green) of size $(2,564 \times 900)$ pixels, the first dimension being the number of time samples and the second being the size of the pushbroom camera; we will keep this convention thereafter. The acquisition rate of 1D images from the pushbroom sensors is 770 Hz, thus the 2D images are roughly recorded in 3.3 seconds. Those datasets were simulated in a way which recreates real-life acquisition conditions with pushbroom sensor noise, ground elevation of the observed scene, radiometric distortions, and realistic attitude variations of the satellite. We did not have access to the simulator as being the property of EADS Astrium, however this data is considered as difficult as could be real data.

A dataset is composed of the warped images, the attitude variations of the satellite which have disturbed the acquisition process and are the groundtruth for our experiments, the focal plane configuration which allows to design the kernel of Equation (10) for a pair of images, and the numerical derivatives of the warp function W with respect to x and t . These latter derivatives are computed using a finite differences method, knowing the position of the satellite and

TABLE 1
Estimated Hyperparameters for the Different Dataset with the Evidence Procedure of Section 5.1

	M3, M4 and M5					M1 and M2		
	σ_I	σ_{a_s}	$\sigma_{a_{nt}}$	σ_{b_s}	$\sigma_{b_{nt}}$	σ_{θ}	σ_I	σ_{θ}
init.	0.05	0.005	0.05	0.005	0.05	0.03	0.05	0.03
D1	0.016	$7 \cdot 10^{-4}$	0.044	$4 \cdot 10^{-4}$	0.19	0.11	0.011	0.13
D2	0.020	$8 \cdot 10^{-4}$	0.087	$4 \cdot 10^{-4}$	0.21	0.05	0.014	0.07
D3	0.015	$4 \cdot 10^{-4}$	0.076	$3 \cdot 10^{-4}$	0.19	0.05	0.013	0.08
D4	0.012	$1 \cdot 10^{-4}$	0.006	$1 \cdot 10^{-4}$	0.01	0.08	0.010	0.10

The first line contains the initial values before the optimization process.

with a digital elevation model of the ground. Specificities of the dataset are:

- Dataset 1 ($D1$), Dataset 2 ($D2$), and Dataset 3 ($D3$) contain multispectral images with respectively high frequency, low frequency, and again low frequency attitude variations.
- Dataset 4 ($D4$) is a tricky case for which images are of the same radiometric modality. However it will be treated as a multispectral set to see how our method performs compared to $M1$ which should be the most appropriate one. It contains low and middle frequency attitude variations.

For all datasets, the pushbroom cameras location on the focal plane is $[1.5, 35, 75, 95]$ in pixel unit for respectively the panchromatic, blue, green, and red channels and according to Fig. 3. As an example, δ equals 20 between the green and red cameras, and the discrete kernel k^{δ} for this pair of images equals $[-10 \dots 01]^T$, with 19 zeros in it.

6.3 Estimation of Hyperparameters

We used the Matlab optimization function `fminunc` to solve Equation (33) with the analytical gradient provided. The most time consuming operation in this expression is to compute the determinant of matrix G , as it is of size $(\tau + 2n_p) \times (\tau + 2n_p)$. To reduce this time to a few hours, we trained the hyperparameters on 10 randomly chosen image patches of size (140×30) for each dataset. We kept the same hyperparameters for all pairs of images in a dataset, which has not induced any significant loss of accuracy in our experiments. Table 1 gives the estimated hyperparameters.

Those same values were repeatedly obtained several times with the same initialization and other image patches. As the marginal likelihood is a function which has multiple local minima as observed in [35], the initial values may change the results. It is however hard to quantify the sensitivity of the method in this multidimensional problem. The initial hyperparameters were manually chosen without knowing the groundtruth. There is no guarantee that a manually tuned model cannot perform better, but one would need the groundtruth to drive the hyperparameters selection, which is not available in a real case scenario.

Overall, the learned hyperparameters for the radiometric coefficients ($\sigma_{a_s}, \sigma_{a_{nt}}, \sigma_{b_s}, \sigma_{b_{nt}}$) are pretty much the same for datasets $D1$, $D2$, and $D3$. One can expect to observe a similar relationship between the panchromatic band and the other channels almost whatever the observed scene is. Dataset

D_4 , which is not a multispectral set, shows different results for the hyperparameters values: it enforces spatial smoothness of $\mathbf{a}$ and $\mathbf{b}$ (smaller values for σ_{a_s} and σ_{b_s}) and keeps their values close to 0 and 1 respectively (again smaller values for $\sigma_{a_{nt}}$ and $\sigma_{b_{nt}}$). The image registration model tends to Equation (9), which is in agreement with a unimodal set of images. Slight differences are observed on the estimated noise parameter of the images σ_I and on the prior parameter of the attitude variations σ_θ which indeed depends on the frequency content of θ . It should be noted that we chose to keep the same σ_θ value for all angles of the attitude because the satellite is likely to oscillate in the same manner around its three rotation axis.

6.4 Attitude Estimation Results

We present attitude estimation results on the roll and the pitch angles for the previously described datasets. No result on the yaw is given as it causes deformations lower than two hundredth of a pixel on our images, it is too small to be recovered with such a complex warp and in a multimodal image registration case. Notice that the literature shows at best pixel accuracy of one hundredth of a pixel in unimodal and five hundredth of a pixel in multimodal cases for translational warps [9], [36].

The evaluation criterion is the standard deviation of the error with respect to the time, denoted ϵ , between the real attitude, denoted $\tilde{\theta}_t$, and the estimated one. This gives:

$$\epsilon = \sqrt{\sum_{t=1}^{\tau} (\tilde{\theta}_t - \theta_t)^2}, \quad (41)$$

and averaged over the roll and the pitch. We chose the Matlab environment to code the algorithms and process the data. Except for the hyperparameter learning procedure which needs the optimization toolbox, no specific function was used and all other algorithms were manually implemented on a middle end desktop computer (core i5 with 8 GB of RAM). To reduce processing time, images were cut into five chunks of 512 pixels along the time axis and 300 pixels along the pushbroom camera axis. Let us remark that it is not necessary to consider the full acquisition of the pushbroom camera, which is 900 pixels here, to estimate the attitude variations θ_t at a given time t . Except for optical flow purposes, pixel based image registration methods usually do not need to use all the pixels of the images to estimate a parametric warp [2]. The location of the (512×300) chunks along the pushbroom camera axis is randomly chosen. Results are given in Table 2; the complete processing time is the result of registering the five chunks of images one after another. As one might expect, method $M1$ gives the best computational time while being quite accurate; if speed is a concern, one should use this technique to register its images. On the other hand if accuracy is crucial, which is often the case with satellite images, methods $M3$, $M4$, and $M5$ offer the best alternatives, with $M4$ offering the best compromise between speed and accuracy. Most surprising results probably come from method $M5$, which is as good as $M3$ but without the need to estimate any radiometric coefficients during the image registration process. However, this method needs a more efficient procedure to be solved, as it

TABLE 2
Attitude Estimation Results on Datasets for Each Method with Respect to Criterion (41)

	D1		D2		D3		D4	
	ϵ	time	ϵ	time	ϵ	time	ϵ	time
M1	0.072	27	0.043	26	0.091	24	0.056	23
M2	0.128	63	0.121	75	0.067	42	0.101	59
M3	0.036	1,215	0.024	1,255	0.033	1,095	0.056	1,199
M4	0.037	281	0.023	254	0.033	245	0.057	261
M5	0.040	1,821	0.024	2,016	0.034	2,675	0.061	2,434
M6	0.055	477	0.044	263	0.071	256	0.060	419

The error ϵ is in pixel unit and the time is in seconds. Best results for a given dataset are in bold face.

is by far the most time consuming one. Poor results of method $M2$ may be explained by the high pass filtering technique which reveals edges on the image but also increases the noise level, even though the DoG operator allows to tune the amount of smoothing in the filter. This solution is more adapted to problems where the modalities between the images are strongly different, like in [18] or [37]. Again in this application, the panchromatic, the red, the blue, and the green images are already well correlated such that discarding the radiometric properties in the warp estimator is still fine. The CVPR method $M6$ gives mixed results; if a more informative prior on the attitude variations may improve the registration process, radiometric corrections are undoubtedly better to consider. To its credit, the autoregressive model usually needs long term signals to efficiently learn its coefficients, which could not be the case here on the small chunks of 512 time samples.

Fig. 5 shows images registration and attitude estimates results for datasets $D3$ and $D4$; they were obtained with method $M4$ on large portions of images ($2,400 \times 240$ pixels) in few minutes. One can notice two different frequency contents between both datasets, the amplitude of the angles which is of several microradians, the corresponding amplitude of the warps on the images as given by the thin white strip which delimits the size of a pixel, and the error between the estimate and the groundtruth.

Fig. 6 presents a (130×130) color image patch from $D1$ before and after the registration process. Bleeding artifacts are visible on the edges of the left image (especially on the white shape) due to the attitude variations, those are corrected on the right image.

Fig. 7 summarizes the benefits of taking into account the radiometric differences between two multispectral images. Images $\mathbf{i}$ and $\mathbf{j}_\theta$, respectively from the panchromatic and the red channels of $D1$, look pretty similar but the difference $|\mathbf{i} - \mathbf{j}_\theta|$ clearly brings out their different nature. After radiometric rectification with $\mathbf{a}$ and $\mathbf{b}$, the statistics of the error $|\mathbf{a} + \mathbf{b}\mathbf{i} - \mathbf{j}_\theta|$ are much closer to a random Gaussian noise.

6.5 Statistics of Estimators

Given Table 2, one may ask whether those results could be obtained repeatedly and with fewer pixels. To this task, we randomly extracted patches of images from dataset $D1$ of size $(512 \times n_t)$, where n_t is a varying number of pixels along the pushbroom camera axis, and ran all the methods




Fig. 5. Results for datasets *D3* (a) and *D4* (b) using method *M4*; on the top is the warped image acquired by the satellite (panchromatic in both cases). The error images are the pixels intensity subtraction between the undistorted panchromatic image (the one the satellite should have recorded if it was not oscillating; it is the groundtruth) and respectively: the acquired panchromatic image (top), and the rectified panchromatic image (bottom) after attitude estimation. Bottom plots display estimated roll and pitch in black curves with the error with respect to the real attitude in gray curves; the thin white strip delimits the size of a pixel.




Fig. 6. Color image patch from dataset *D1* before and after the registration process. Several bleeding artifacts occur on the edges of the left image due to the attitude variations. Those are corrected after the registration process on the right image.

100 times for a given n_l to get an average error $\bar{\epsilon}$ and a standard deviation σ_ϵ of the criterion (41) for each estimator. Results are visible in Fig. 8.

One can see that method *M3* has both the best performances in terms of mean and variance of the error, followed closely by *M4*. The approximation scheme in the two step algorithm does not decrease the performances of the estimator when $n_l > 15$ while being four times faster than the global algorithm.

7 DISCUSSION AND EXTENSIONS

We have presented a complete framework for attitude variations estimation of a satellite using the images it has recorded with its pushbroom sensors. It makes use of the acquisition process geometry, it automatically deals with radiometric differences between the multi-spectral images and insures a high accuracy estimation of the warps. It efficiently solves the problem in a




Fig. 7. Comparison between two image patches *i* and j_θ , of size (100×100) pixels from the panchromatic and the red channels of *D1*, after the registration process of method *M3*. The error image $|a + bi - j_\theta|$ is much closer to a random Gaussian noise than $|i - j_\theta|$.




Fig. 8. Performances of the methods with respect to the number of horizontal pixels n_l in the images. Results were averaged over 100 runs on randomly chosen image patches from dataset *D1* at a given size of $(512 \times n_l)$ pixels.

Bayesian formalism which allows to: use priors for the attitude variations and the radiometric coefficients, automatically learn the hyperparameters of the model, and marginalize the radiometric coefficients while keeping a good accuracy. The experimental part motivates the choice of the method *M4*.

The main challenge with satellite datasets is to efficiently use the large amount of data which is composed of multiple gigapixel images. For instance, instead of aligning all the images of the focal plane, one could choose the best pair of images for which the corresponding kernel spectrum, as given in Fig. 4, includes the frequencies of the attitude variations. Another solution could be to select the image patches with the appropriate number of pixels n_l along the camera axis to estimate the attitude variations with a good confidence value. However to our best knowledge, there is no clear answer to this question in the pixel based image registration literature, though this was a great concern for feature based techniques like in [38]. The main difficulty is that the estimator performances is linked to the image content, which is somehow intuitive but greatly demonstrated in [36] using the Cramer Rao bound. There should be a link between the image content and the number of pixels n_l to use in the registration process which insures an estimator with a sufficiently small variance.

Another future research perspective could be to better understand the warp function $\mathcal{W}$, as it is currently numerically approximated. Given an analytical model for this function and the amount of data, one may be able to jointly estimate the attitude variations of the satellite as well as the elevation map of the scene.

ACKNOWLEDGMENTS

This work was fully sponsored by EADS Astrium, the European satellite manufacturer, in collaboration with INRIA Grenoble.

REFERENCES

- [1] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control*, vol. 19, no. 6, pp. 716–723, Dec. 1974.
- [2] S. Baker and I. Matthews, "Lucas-Kanade 20 years on: A unifying framework," *Int. J. Comput. Vis.*, vol. 56, no. 1, pp. 221–255, 2004.
- [3] F. Brunet, A. Bartoli, N. Navab, and R. Malgouyres, "Pixel-based hyperparameter selection for feature-based image registration," in *Proc. Int. Workshop Vis., Model., Vis.*, 2010, pp. 33–40.
- [4] F. J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 6, pp. 679–698, Nov. 1986.
- [5] M. Cramer, D. Stallmann, and N. Haala, "Direct georeferencing using GPS/inertial exterior orientations for photogrammetric applications," in *Proc. Int. Archives Photogrammetry Remote Sens.*, 2000, pp. 198–205.
- [6] F. de Lussy, D. Greslou, and L. G. Colzy, "Process line for geometrical image correction of disruptive microvibrations," in *Proc. Int. Soc. Photogrammetry Remote Sens.*, 2008, pp. 27–33.
- [7] N. Dowson, and R. Bowden, "Mutual information for Lucas-Kanade tracking (MILK): An inverse compositional formulation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 1, pp. 180–185, Jan. 2008.
- [8] J. Drareni, P. Sturm, and S. Roy, "Plane-based calibration for linear cameras," in *Proc. 8th Workshop Omnidirectional Vis., Camera Netw. Non-Classical Cameras*, 2008, pp. 146–156.
- [9] W. Förstner, "Quality assessment of object location and point transfer using digital image correlation techniques," in *Proc. Int. Archives Photogrammetry*, 1984, pp. 197–219.
- [10] P. Gordon, "Airborne pushbroom line scanners: An alternative to digital frame scanners," *Geoinformatics*, vol. 8, no. 1, pp. 50–57, 2005.
- [11] A. Gruen and L. Zhang, "Sensor modeling for aerial triangulation with three-line scanner (TLS) imagery," *Photogrammetrie, Fernerkundung, Geoinformation*, vol. 1, no. 2, pp. 85–98, 2003.
- [12] A. Guimond, A. Roche, N. Ayache, and J. Meunier, "Multimodal brain warping using the demons algorithm and adaptive intensity corrections," *IEEE Trans. Med. Imag.*, vol. 20, no. 1, pp. 58–69, Jan. 2001.
- [13] R. Gupta and R. I. Hartley, "Camera estimation for orbiting pushbrooms," in *Proc. 2nd Asian Conf. Comput. Vis.*, 1995.
- [14] R. Gupta and R. I. Hartley, "Linear pushbroom cameras," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 9, pp. 963–975, Sep. 1997.
- [15] C. Harris and M. Stephens, "A combined corner and edge detector," in *Proc. 4th Alvey Vis. Conf.*, 1988, pp. 147–151.
- [16] H. Hirschmüller, F. Scholten, and G. Hirzinger, "Stereo vision based reconstruction of huge urban areas from an airborne pushbroom camera (HRSC)," in *Proc. 27th DAGM Conf. Pattern Recog.*, 2005, pp. 58–66.
- [17] J. Inglada and A. Giros, "On the possibility of automatic multisensor image registration," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 10, pp. 2104–2120, Oct. 2004.
- [18] M. Irani and I. P. Anandan, "Robust multi-sensor image alignment," in *Proc. Int. Conf. Comput. Vis.*, 1998, pp. 959–966.
- [19] M. Irani and P. Anandan, "About direct methods," in *Proc. Int. Workshop Vis. Algorithms*, 1999, pp. 267–277.
- [20] A. Iwasaki, "Detection and estimation satellite attitude jitter using remote sensing imagery," in *Advances in Spacecraft Technologies*. Rijeka, Croatia: InTech, 2011, pp. 259–272.
- [21] A. Jalobeanu, "Predicting spatial uncertainties in stereo photogrammetry: Achievements and intrinsic limitations," in *Proc. 7th Int. Symp. Spatial Data Quality*, Oct. 2011.
- [22] A. Jalobeanu, L. Blanc-Féraud, and J. Zerubia, "Hyperparameter estimation for satellite image restoration using a MCMC maximum likelihood method," *Pattern Recog.*, vol. 35, no. 2, pp. 341–352, 2002.
- [23] J. Lamard, C. Gaudin-Delrieu, D. Valentini, C. Renard, T. Tournier, and J. Laherrere, "Design of the high resolution optical instrument for the PLEIADES HR earth observation satellites," in *Proc. 5th Int. Conf. Space Opt.*, 2004, pp. 149–156.
- [24] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, pp. 91–110, 2004.
- [25] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. Int. Joint Conf. Artif. Intell.*, 1981, pp. 674–679.
- [26] D. J. C. Mackay, "Bayesian interpolation," *Neural Comput.*, vol. 4, no. 3, pp. 415–447, 1992.
- [27] D. J. C. Mackay, *Information Theory, Inference and Learning Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2002.
- [28] L. Magerand and A. Bartoli, "A generic rolling shutter camera model and its application to dynamical pose estimation," in *Proc. Int. Symp. 3D Data Process., Vis. Transmiss.*, 2010.
- [29] J. Makhoul, "Linear prediction: A tutorial review," *Proc. IEEE*, vol. 63, no. 4, pp. 561–580, Apr. 1975.
- [30] R. Perrier, E. Arnaud, P. Sturm, and M. Ortner, "Estimating satellite attitude from pushbroom sensors," in *Proc. IEEE 23th Conf. Comput. Vis. Pattern Recog.*, 2010, pp. 591–598.
- [31] R. Perrier, E. Arnaud, P. Sturm, and M. Ortner, "Sensor measurements and image registration fusion to retrieve variations of satellite attitude," in *Proc. 10th Asian Conf. Comput. Vis.*, 2010, pp. 361–372.
- [32] D. Poli, "General model for airborne and spaceborne linear array sensors," in *Proc. Int. Archives Photogrammetry Remote Sens.*, 2002, pp. 177–182.
- [33] O. Punska, "Bayesian approaches to multi-sensor data fusion," Ph.D. dissertation, Dept. Eng., Cambridge, U.K., Cambridge Univ., 1999.
- [34] P. Rademacher and G. Bishop, "Multiple-center-of-projection images," in *Proc. 25th Annu. Conf. Comput. Graph. Interactive Techn.*, 1998, pp. 199–206.
- [35] C. E. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [36] D. Robinson, and P. Milanfar, "Fundamental performance limits in image registration," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1185–1199, Sep. 2003.
- [37] A. Roche, G. Malandain, N. Ayache, and S. Prima, "Towards a better comprehension of similarity measures used in medical image registration," in *Medical Image Computing and Computer Assisted Intervention*. New York, NY, USA: Springer, 1999, vol. 1679, pp. 555–566.
- [38] J. Shi and C. Tomasi, "Good features to track," in *Proc. Comput. Vis. Pattern Recog.*, 1994, pp. 593–600.
- [39] R. Szeliski, "Image alignment and stitching: A tutorial," *Found. Trends Comput. Graph. Vis.*, vol. 2, no. 1, pp. 1–104, 2006.
- [40] R. Szeliski, *Computer Vision: Algorithms and Applications*. New York, NY, USA: Springer, 2010.
- [41] B. Zitova and J. Flusser, "Image registration methods: A survey," *Image Vis. Comput.*, vol. 21, no. 11, pp. 977–1000, 2003.
- [42] A. Zomet and S. Peleg, "Multi-sensor super-resolution," in *Proc. IEEE 6th Workshop Appl. Comput. Vis.*, 2002, p. 27.



Régis Perrier received the master's degree from South Toulon-Var University, France and the PhD degree from Grenoble University, France, in 2005 and 2011, respectively. As an engineer, he has work on several topics related to signal processing for seismic event localization, and focused during the PhD degree on computer vision and remote sensing. He was a lecturer in computer science for Grenoble University in 2011. He is currently a research engineer at the CEA-Leti on inertial and magnetic sensors based motion capture. His research interests include sensor fusion, bayesian filtering, computer vision, and machine learning.



Elise Arnaud received the PhD degree in signal processing and telecommunications from the University of Rennes, France, in 2004. Since September 2006, she has been an assistant professor at Université Joseph Fourier, Grenoble, France. She started her professional career with a postdoctoral position at the University of Buenos Aires, Argentina, followed by a one-year research at the University of Genova, Italy. Her main research interests include statistical models for analysing dynamic scenes from image sequences. She is currently working on designing mathematical and computer tools to help decision-maker for a transition to sustainability at a regional scale.



Peter Sturm received the MSc degrees from the National Polytechnical Institute of Grenoble (INPG), France and the University of Karlsruhe, Germany, both in 1994, and the PhD degree from INPG in 1997, with Long Quan as an advisor. His PhD thesis was awarded the SPECIF Award (given to one French PhD thesis in Computer Science per year). After a two-year postdoc at Reading University, working with Steve Maybank, he joined INRIA on a permanent research position as chargé de recherche in 1999. Since

2006, he has been the directeur de recherche (the INRIA equivalent of professor). In 2009-2010 he spent a one-year sabbatical at CAMP, TU Munich, Germany.



Mathias Ortner received the BSc degree in aeronautics and space sciences from the French School of Aeronautics and Space and the master's degree in applied mathematics from the University of Toulouse, France, in 2001. He received the PhD degree in computer science and image processing from the University of Nice-Sophia Antipolis, France, in 2004. From 2004 to 2006, he was a postdoctoral research visitor in Professor Nehorais team, first at the University of Illinois, Chicago, and then at Washington University in Saint Louis. He is currently a research engineer for Astrium, an EADS company, in Toulouse. His research interests include stochastic methods and Bayesian models for image processing.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**