

Making 3D movies of Northern Lights

Eric Hivon^{1,*}, Jean Mouette¹ and Thierry Legault^{2,a}

¹ Sorbonne Universités, UPMC Univ. Paris 6 & CNRS (UMR 7095), Institut d'Astrophysique de Paris, 98 bis Bd Arago, 75014 Paris, France

² THALES, 19 Avenue Morane Saulnier, 78140 Vélizy Villacoublay, France

Received 4 June 2017 / Accepted 16 August 2017

Abstract – We describe the steps necessary to create three-dimensional (3D) movies of Northern Lights or Aurorae Borealis out of real-time images taken with two distant high-resolution fish-eye cameras. Astrometric reconstruction of the visible stars is used to model the optical mapping of each camera and correct for it in order to properly align the two sets of images. Examples of the resulting movies can be seen at <http://www.iap.fr/aurora3d>.

Keywords: Aurora / visualisation / technical / image processing / algorithm

1 Introduction

Aurorae borealis have always fascinated humans, who have long tried to report their observations by the best means available at the time. If the earliest known record of an Aurora was written on a clay tablet 567 BC in Babylon (Stephenson *et al.*, 2004), they have since been filmed in color with specially designed high sensitivity cameras¹ in the mid-1980s (Hallinan *et al.*, 1985), and filmed in 3D for the first time in 2008.²

In this paper we detail the steps taken to generate 3D movies of Aurorae Borealis shot with high-resolution wide field cameras, generating movies fit for projection on hemispheric planetariums. To our knowledge, these 3D movies are the first ones showing Northern Lights ever produced thanks to post-shooting astrometric reconstruction. Since they are produced from real-time movie images instead of the usual time-lapsed sequences, they can be used for education and public outreach and give a glimpse at the rich and rapidly moving three-dimensional structure of the Auroras. While they are close to the real-world sensations of Aurora watching in terms of colors, saturation, speed of evolution, and sheer size on the sky (if seen in a planetarium), such movies add the extra information of the stereoscopy, which is inaccessible to a single human observer on the field. Now that such movies are made available, they will hopefully

prompt the interest of scientists studying the Aurorae properties, who would like to use them, or build upon them, in their work.

The image taking procedure and camera synchronization is described in Section 2. The image processing steps allowing a good rendition of the stereoscopy is described in Section 3, while Section 4 shows how such 3D images and movies can be viewed. Section 5 is devoted to a conclusion and perspectives.

2 Observational setup

Observations are made from Norway at the approximate latitudes and longitudes of 69.3°N and 20.3°E, with two cameras separated by either 6 or 9 km, depending on the local weather conditions. Since the Aurorae mostly happen in the upper atmosphere at ~100 km above the ground, their parallax is expected to be large enough between the two cameras to construct their 3D rendering via binocular vision. As much as possible, the cameras point to the zenith, and are then rotated around the vertical so that the Northern Star is in the lower part of the image, aligning the long axis of the image with the West-East direction, as shown in Figure 1 for one of the cameras. Their positions are measured by GPS receivers. In what follows, the western and eastern cameras are, respectively, dubbed left and right.

2.1 The cameras

Each of the two cameras used is a Sony $\alpha 7s$, on which is mounted a Canon Fisheye lens of focal length $f = 15$ mm with an aperture of $f/2.8$. In order to further enlarge the field of view, and allow the use of a Canon lens on the Sony E-mount camera, it is coupled to a Metabones Speed Booster that

^a Astrophotography, <http://www.astrophoto.fr>.

*hivon@iap.fr

¹ <http://auroraalive.com/multimedia/autoformat/getswf.php?videoSite=aurora&videoFile=aaauroraontv.swf++&videoTitle=Aurora+on+TV>.

² <https://www.newscientist.com/article/dn15147-northern-lights-captured-in-3d-for-the-first-time/>.




Fig. 1. The observation is done from colatitude θ_0 and longitude ϕ_0 , in the Earth-bound xyz referential. A camera-bound referential is defined, with the Z axis along the camera boresight (roughly towards the zenith), the Y axis along the shorter dimension of the camera image, approximately pointing North, and the X axis along the longer dimension of the camera image, pointing East.

reduces the focal length by a factor of 0.71. The camera sensor is a full-frame CMOS 4240×2832 pixel RGBG Bayer matrix,³ with half of the pixels being sensitive to green and the other half spread evenly between red and blue. ISO is set to 12 800, representing a good compromise between noise and sensitivity. A video signal is produced at 25 or 30 frames per second (fps), in Ultra High Definition format (3840×2160 pixels, ratio 16:9), and recorded on an external video recorder, with Apple ProRes HQ codec (approximately 800 Mbits/s). As we will see in Section 3, the field of view is $220^\circ \times 110^\circ$, covering 48% of the sphere, with a resolution at the center of the image of $2.94''/\text{pixel}$.

2.2 Cameras synchronization

In order to achieve a good stereoscopy out of the movies produced, we first have to make sure that their images are properly synchronized during their processing and visualisation. Keeping, on the field, the two distant cameras and/or their recorders synchronized thanks to a common absolute GPS-based time signal would have been ideal, but out of reach of this project. Instead, to achieve a good relative synchronization of the two movies, we first made sure that the internal clocks of the two recorders agreed to one second or less before starting every observing nights; we then flashed a light in front of the two cameras put side by side and connected to their turned-on recorder before setting up the cameras at their respective distant location. This gives us a good estimate of the offset and possible drift in time between the internal clocks of the two recorders as they write a time-code in the images metadata.

³ <http://www.imaging-resource.com/PRODS/sony-a7s/sony-a7sDAT.HTM>.

Finally, in post-production, but before starting the 3D reconstruction pipeline, we look in the filmed sequences for the occurrence of bright, thin and fast moving Auroras structures, and compare closely their evolution from frame to frame between the two cameras, in order to find the best matching pair of frames. This provides the residual time offset between the two series of recorded time-stamps, and allows us to re-synchronize the two movies at a level compatible with the image rate (25 or 30 fps), assuming the relative time drift to be negligible over the duration of a sequence (a few minutes).

3 Image processing

Time-lapsed stereoscopic observations of Aurorae are described for wide field cameras in Kataoka *et al.* (2013), and for fast narrow field cameras in Kataoka *et al.* (2016), following in either cases the procedure described in Mori *et al.* (2013). We are implementing a similar astrometric method: since the actual pointing, orientation and relative position of the two cameras are only known with limited accuracy, and because each camera can distort the image differently, the positional astrometry of bright stars identified in the images is used to determine and characterize the cameras geometrical settings and optical responses, in order to properly realign the left and right set of images.

3.1 Finding the stars

The 5th edition of the Kitt Peak National Observatory catalog,⁴ adapted from the Yale catalog,⁵ lists the J_{2000} equatorial coordinates and magnitude of more than 9000 stars of magnitude ≤ 6.5 . After setting correctly the minus sign on the magnitude of its four brightest stars, it was used as a catalog of bright stars potentially seen in the left and right images.

Candidate stars in the images are identified as follows. Since the Aurorae filmed here are mostly green in color, a red channel frame is expected to be less exposed to them than the other channels, and is used to detect stars. We checked that using the green channel, expected to have a better intrinsic resolution, did not change much the final result, while it increased slightly the risk of false detection. The chosen image is convolved (via Fast Fourier Transform operations) with a difference of Gaussians (DoG) filter, defined in pixel space as

$$F(x, y) = \frac{\exp\left(-\frac{x^2+y^2}{2\sigma_1^2}\right)}{2\pi\sigma_1^2} - \frac{\exp\left(-\frac{x^2+y^2}{2\sigma_2^2}\right)}{2\pi\sigma_2^2}, \quad (1)$$

where x, y are integer pixel indices, $\sigma_1 = 1$ and $\sigma_2 = 1.6$, and one then looks for local maxima in the filtered image. The value of $\sigma_2/\sigma_1 = 1.6$ picked here is generally recommended as a good approximation of the Mexican hat or Marr wavelet (Marr and Hildreth, 1980), often used to look for point sources in astrophysics (Coupinot *et al.*, 1992), but values as large as $\sigma_2/\sigma_1 = 3$ gave almost identical results.

⁴ <http://www-kpno.kpno.noao.edu/Info/Caches/Catalogs/BSC5/catalog5.html>.

⁵ <http://tdc-www.harvard.edu/catalogs/bsc5.html>.

For practical reasons, we keep our cameras in focus during the totally of the shooting sequences. As a consequence, the Airy diameter of the optical system ($\sim 3 \mu\text{m}$) is smaller than the pixels physical size of the CMOS sensor ($8.40 \mu\text{m}$), and the observed stars generally would not be resolved. We tried three different techniques to determine the position of a star on the frame: (a) that of the locally brightest pixel of the DoG filtered map, or the centroid of the filtered map on a (b) 3×3 or (c) 5×5 patch of pixels centered on that same brightest pixel. For reasons that will be detailed in [Section 3.3](#), the option (b) which provides sub-pixel accuracy was preferred. Since it is difficult to estimate exactly the error made on determining the star position, especially since we start from a (red) image interpolated with proprietary algorithms from the RGBG detector matrix, we will simply assume the error in position in (b) to be the same as in the discretization scheme (a) and determined by the nominal pixel size of the image, as described in [Appendix B](#). Note that over-estimating by a factor 10 the error bar on each star position would not change the numerical value of any parameter in the multi-dimensional fit being done, it will only increase by a factor 10 the error bar associated to each parameter.

The $N_b \simeq 40$ brightest sources found by this procedure are compared to the $N_s \simeq 40$ brightest stars in the Northern hemisphere, as listed in the input catalog described above.

3.2 Modeling the projection

3.2.1 3D rotation of the sky

At a time t , a star (j) having the colatitude⁶ and longitude $(\delta_E^{(j)}, \alpha_E^{(j)})$ in equatorial coordinates, will have the angular coordinates $(\delta_c^{(j)}, \alpha_c^{(j)})$ in a camera-bound referential, such as the one illustrated in [Figure 1](#), so that the respective 3D coordinate vectors are

$$\mathbf{r}_i^{(j)} \equiv \begin{pmatrix} x_i^{(j)} \\ y_i^{(j)} \\ z_i^{(j)} \end{pmatrix} = \begin{pmatrix} \sin \delta_i^{(j)} \cos \alpha_i^{(j)} \\ \sin \delta_i^{(j)} \sin \alpha_i^{(j)} \\ \cos \delta_i^{(j)} \end{pmatrix}, \quad (2a)$$

for $i = E$ or c . This change of coordinates is a rotation parameterized by the three (time dependent) Euler angles ψ, θ, φ

$$\mathbf{r}_c^{(j)} = \mathbf{R}(\psi, \theta, \varphi) \mathbf{r}_E^{(j)}, \quad (2b)$$

with

$$R(\psi, \theta, \varphi) \equiv \begin{pmatrix} \cos \psi & -\sin \psi & 0 \\ \sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix} \cdot \begin{pmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (2c)$$

The solid-body rotation of the sky described in equation (2) is an approximation that ignores subtle distortions like the atmospheric refraction, which will be discussed later on, or the relativistic aberration due to the yearly Earth motion around the Sun. The latter is fairly small, with an apparent displacement of the source varying between 0 and $\sim 20.5''$

across the sky, and will affect almost identically the two cameras since they observe the same area of the sky. It should therefore not affect the relative alignment of the images provided by the two cameras. For a camera located at a time t at the Earth-bound colatitude θ_0 and longitude ϕ_0 , pointing exactly to the zenith and rotated by ψ_0 around the vertical (with $\psi_0 = 0$ in the current configuration), then, ignoring the Earth nutation,

$$\psi = -90 - \psi_0, \quad (3a)$$

$$\theta = -\theta_0, \quad (3b)$$

$$\varphi = -(\phi_0 + 15S(t) \bmod 360), \quad (3c)$$

where $S(t)$ is the Greenwich Mean Sidereal Time of observation, expressed in hours, which parameterizes the rotation of the Earth with respect to the distant stars:

$$S(t) = 18.697374558 + 24.06570982441933 \times (J(t) - J_{2000}), \quad (4)$$

with $J(t)$ the Julian day of observation and $J_{2000} = 2\,451\,545$ is the Julian day starting at 12:00:00 UTC on January 1, 2000 ([Meeus 1991](#)). However, as noted previously, the limited accuracy on the actual shooting time, position and orientation of the cameras will impact these parameters. The values recovered for (ψ, θ, φ) , as described in [Section 3.3](#), are well within 1° of those expected from equation (3), with the larger discrepancy affecting the harder to set-up azimuthal angle ψ .

3.2.2 Radial distortion of the optics

A source located at an angular distance δ_c of the camera bore-sight will appear on the imaging CMOS sensor at a distance r of the focal point, as illustrated in [Figure 2](#) and we model this distance as

$$r^{(j)} = f_1 r_{ES}(\delta_c^{(j)}) + f_a r_{ES}^a(\delta_c^{(j)}) + f_b r_{ES}^b(\delta_c^{(j)}), \quad (5a)$$

where

$$r_{ES}(\delta_c) \equiv 2\sin(\delta_c/2), \quad (5b)$$

is the equisolid radius ideally expected for the fisheye and reducer lenses being used, f_1 is the combined focal length of the fisheye and focal reducer, and f_a and f_b model any non-linear departure from the expected projection, for the integer numbers a and b . We found the combination $(a, b) = (3, 5)$ to be slightly better than $(2, 3)$, by requiring smaller amplitude of the corrections, i.e., $f_3/f_1 \simeq -0.008$ and $f_5/f_1 \simeq 0.004$, instead of $f_2/f_1 \simeq -0.02$ and $f_3/f_1 \simeq 0.01$, in the range probed by the stars ($\delta_c < 90^\circ$ and $r_{ES} < 1.41$), and by returning slightly smaller residual in the comparison of the model with the observations.

Since the optics is pointing toward the zenith, these non-linear terms can result from both radial non-idealities of the optics and, as suggested by [Mori et al. \(2013\)](#), the atmospheric refraction of the incoming light, which increases with the camera-bound colatitude of the source. However, as shown in [Appendix A](#) and in [Figure 2](#), atmospheric refraction is too small in amplitude to explain the distortions seen, since we would then get $f_3/f_1 = 2f_5/f_1 \simeq -8.24 \times 10^{-5}$.




Fig. 2. Radial mapping of the camera, relating the camera-bound colatitude δ_c of a sky object, and its projected radial distance r on the camera image, measured in pixels. In the case of stars (red diamonds) the measured radial distance comes with the discretization error bar of size $\pm 1/\sqrt{12} \approx \pm 0.2887$ (see Appendix B) more visible on the lower panel. The black dots show a simple equidistant mapping $r \propto \delta_c$. The magenta dashes show the equisolid projection of equation (5b) expected for the fisheye and focal reducer lenses used here; the green solid line show the best fit distorted radial model of equation (5a); while the blue dotted-dashed line illustrate the impact of atmospheric refraction on observations with a perfect equisolid optical system (see Appendix A). The upper panel shows the raw radial distance r , while the lower panel shows the difference between r and the equisolid model.

3.2.3 2D offset of the camera sensor

Finally, the geometric center of the camera sensor, identified as $(x, y) = (0, 0)$ may not match exactly the optics focal point, and we allow for the shifts Δ_x, Δ_y , so that a source j will appear on the sensor at the location

$$x^{(j)} = r^{(j)} \cos \alpha_c^{(j)} + \Delta_x, \quad (6a)$$

$$y^{(j)} = r^{(j)} \sin \alpha_c^{(j)} + \Delta_y, \quad (6b)$$

measured in units of pixels, from the nominal CMOS sensor center. As listed in Table 1, we find these shifts to be, respectively, $|\Delta_x| \sim 20$ and $|\Delta_y| \sim 5$ pixels, depending on the camera considered, and the conditions of observation.

⁶ In what follows, and unless stated otherwise, the angles will be expressed in degrees; and we will prefer the colatitude δ measured southward from the North pole to the usual latitude δ' measured northward from the equator, the two being related by $\delta + \delta' = 90^\circ$.

3.3 Fitting the parameters

Combining equations (2), (5) and (6), the position of a star on the camera sensor is related to its equatorial coordinates via

$$\begin{pmatrix} x^{(j)} \\ y^{(j)} \end{pmatrix} = D \left(\psi, \theta, \varphi, f_1, f_a, f_b, \Delta_x, \Delta_y; \begin{pmatrix} \delta_E^{(j)} \\ \alpha_E^{(j)} \end{pmatrix} \right), \quad (7)$$

where the eight parameters $(\psi, \theta, \varphi, f_1, f_a, f_b, \Delta_x, \Delta_y)$ are assumed unrelated between the two cameras.

In order to determine these parameters for a camera, we begin by setting the angles (ψ, θ, φ) to the values given in equation (3), and letting all the other parameters to 0, except for f_1 which is set to the nominal focal length of the optical system. Applying equation (7) to the N_s brightest stars of the catalog, we find that for each of a handful of them (N_m), the computed location is within a few pixels of a bright source identified on the image. Assuming the star and the close bright pixel to be the same object, we then look for the set of parameters that minimize the $2N_m$ discrepancies in (x, y) coordinates. To do this non-linear fitting, we tried and compared two different IDL implementations of the Leven-

Table 1. Parameters of the camera response, for frames taken from six different shooting sequences with the left (rows L1–L6) and right (rows R1–R6) cameras. All runs were shot on consecutive nights, except for #2–#4 which were produced on the same night, and #6 obtained four months later. Figures 2 and 3 correspond to row R2. The focal length f_1 was converted from pixels to mm assuming a pixel size of $8.40 \mu\text{m}$. The errors quoted on the parameters $f_1, f_3, f_5, \Delta_x, \Delta_y$ are those returned by the non-linear fitting procedure, based on the assumed discretization error on star location. The two rightmost columns show the mean and maximum discrepancy between the measured and modelled positions of the N_m stars found in the frame.

Camera & run	N_m	f_1 (mm)	$1000f_3/f_1$	$1000f_5/f_1$	Δ_x (pixels)	Δ_y (pixels)	Mean discr. (pixels)	Max discr. (pixels)
L1	22	9.811 ± 0.003	-6.1 ± 0.5	3.8 ± 0.3	-17.7 ± 0.5	6.0 ± 0.4	0.30	0.71
L2	28	9.815 ± 0.003	-7.0 ± 0.4	4.4 ± 0.2	-19.8 ± 0.4	-0.1 ± 0.4	0.29	0.68
L3	27	9.817 ± 0.003	-7.4 ± 0.4	4.5 ± 0.2	-19.0 ± 0.4	-0.0 ± 0.4	0.34	0.66
L4	25	9.819 ± 0.003	-8.2 ± 0.4	5.0 ± 0.3	-21.2 ± 0.4	0.3 ± 0.4	0.31	0.58
L5	22	9.818 ± 0.003	-6.6 ± 0.5	4.0 ± 0.3	-7.7 ± 0.5	2.0 ± 0.3	0.41	1.07
L6	23	9.859 ± 0.004	-5.6 ± 0.6	3.0 ± 0.4	-10.2 ± 0.4	4.3 ± 0.3	0.33	0.98
R1	23	9.890 ± 0.003	-5.8 ± 0.5	3.8 ± 0.3	14.0 ± 0.5	5.0 ± 0.4	0.21	0.42
R2	25	9.891 ± 0.003	-5.3 ± 0.5	2.9 ± 0.3	13.9 ± 0.4	5.3 ± 0.5	0.26	0.58
R3	23	9.893 ± 0.003	-6.7 ± 0.4	4.3 ± 0.3	14.1 ± 0.5	4.8 ± 0.5	0.22	0.44
R4	23	9.885 ± 0.004	-3.7 ± 0.6	1.9 ± 0.5	14.2 ± 0.4	4.9 ± 0.5	0.26	0.55
R5	25	9.897 ± 0.003	-8.3 ± 0.4	5.6 ± 0.2	8.1 ± 0.5	0.4 ± 0.3	0.28	0.53
R6	22	9.888 ± 0.004	-5.3 ± 0.6	3.4 ± 0.3	20.5 ± 0.4	8.8 ± 0.3	0.33	0.85

berg–Marquardt algorithm (Marquardt 1963, Press et al. 1992): the `curvefit` routine included in the standard IDL library, and its drop-in replacement `mpcurvefit`⁸ by Markwardt (2009) based on the MINPACK algorithm of Moré (1978). We found the best fit parameter values to be nearly identical, but only the latter routine returned meaningful error bars on those parameters. Injecting these new parameters in equation (7), the number of coincidences N_m is increased, and we again look for a new estimate of the parameters minimizing the $2N_m$ discrepancies. The process is repeated a few times, rapidly converging toward a stable number of matches ($22 \leq N_m \leq 28$ depending on the image being treated), and providing a stable set of fitted parameters.

Figure 3 compares the measured position of the bright sources and the computed position of their matching stars in the best fit model of the camera response, for one specific image. Their discrepancies, shown as magenta arrows (after multiplication by 200) seem fairly randomly distributed and do not exhibit any clear trend. The residual distances have a mean value of 0.29 (in pixel units), with a worst case of 0.68. The mean residual found is compatible with the error created by assigning an integer value to the star coordinate in the image, which is $\simeq 0.3826$, as shown in Appendix B.

As shown in Table 1, the same reconstruction procedure was applied, for each of the two cameras, on images extracted from six different shooting sequences, filmed on four different nights spread over a four month period. The optics related parameters ($f_1, f_3, f_5, \Delta_x, \Delta_y$) and the level of residual discrepancies in position, were found to be quite stable for each camera, with the largest relative changes affecting Δ_x, Δ_y when the lenses were unmounted and remounted between observations performed on different nights. The relatively large changes of f_3 and f_5 between runs may be due to a partial (anti-) correlation between these two parameters, which also shows

on the fact that the combination $f_3 + f_5$ varies less between runs than either f_3 or f_5 . However, this degeneracy does not hamper the modelling of equation (7) as long as f_3 and f_5 are treated together, and not considered separately. This suggests, however, that any more sophisticated modelling of the optical response, and in particular of the radial distortion, would require either the use of a basis of orthogonal polynomials, or a physically motivated set of parameters.

In each case, the mean residual error in position is compatible or below what is expected from quantization error of the star coordinates (0.3826 pixels), and we note that the right camera seems to perform a bit better than the left one on this respect. This may be due to the slightly lesser sharpness of the images produced by the right camera, maybe due to its optics, which by blurring lightly the stars makes the sub-pixel determination of their position easier.

As mentioned in Section 3.1, different options were considered to determine the stars position on the camera frame. Using the centroid on a 5×5 patch (so-called option (c)) instead of 3×3 used here made no significant difference on the parameter values and did not improve the residual discrepancies compared to those listed in Table 1. As expected, using instead a discrete pixel location (option (a)) lead to different parameter values, with a shift compatible with the quoted error bars, and larger mean and worst case discrepancies (by $\sim 30\%$ and $\sim 20\%$, respectively). However, for the 12 images tested, the sky to pixel mappings of equation (7) resulting from options (a) and (b), respectively, differed by quite less than one pixel in the region of interest, i.e., everywhere above the horizon, making the resulting processed images and movies of Auroras nearly indistinguishable.

The 8-parameter mapping model considered here therefore seems to generate smaller reconstruction errors than the 6-parameter model considered in Mori et al. (2013), which ignores the non-linear radial distortions (f_a and f_b). On the other hand, in the framework of meteors and transient objects detection, Borovička et al. (1995) proposed a model containing 13 (12 independent) parameters for the modelling of all sky

⁷ <https://harrisgeospatial.com/docs/CURVEFIT.html>.

⁸ <http://purl.com/net/mpfit>.




Fig. 3. Comparison of the bright sources location and the position computed with the optical model, for one frame of one of the cameras. The bright pixels detected in the camera image are shown as black circles, while the expected position of the associated stars, computed by equation (7) with the eight camera parameters obtained in Section 3.3, are shown as red stars. The residual discrepancies between the two, multiplied by 200, are shown as magenta arrows. The thick black dots outline the approximate landscape horizon on the image; the black cross close to the image center shows the optics focal point; and the cyan lines represent the equatorial graticule, with a spacing of 10° in (equatorial) colatitude and 45° in longitude. The Northern Star is marked NS, while the “dipper” part of Ursa Major, in the North-East quadrant, is outlined in grey.

cameras equipped with photographic plates featuring a measurement uncertainty of ~ 1 arcmin. With hour-long exposures and ~ 100 stars for calibration, they reached a residual modelling error of ~ 1 arcmin too. We did not investigate how such a model would have performed in our case. Even though adding more parameters is not problematic in itself for the non-linear fitting procedure, it may lead to further degeneracies between parameters, unless they are carefully designed to be orthogonal, especially if the number of constraints (here, the identified stars) is not large enough to correctly probe each of them. Moreover, inspection of Figures 2 and 3 does not suggest the need for more degrees of freedom.

3.4 Aligning the images

Once the parameters are determined for each of the left and right images, with respect to the absolute equatorial coordinates, one can map the two images (or set of images) onto a common projection of our choice, so that the stars match exactly for the left and right “eyes”, while the structures of the Aurorae will be offset according to their parallax. We chose for instance, for flat screen viewing, and for Figures 4 and 5, to deproject the right image by inverting equation (7) for the right camera parameters, and projected it again by applying that same equation with the left camera parameters. Other choices are possible, such as mapping the two images onto a common equidistant projection, as required by the Dome Master format⁹ used in planetariums.

Apart from the angle φ , which includes the apparent motion of the sky due to Earth rotation, the parameters are not expected to vary during the typical duration of a filmed sequence (2–3 min). We chose to determine these two sets of parameters on the first image of the left and right movies, respectively, and apply the mappings based on them, one frame at a time, to the whole span of each of the two movies, meaning that the stars will appear to slowly rotate on the sky, as they would for a human observer with a very wide field of view.

An extra twist is that the two cameras do not sit on the same latitude, even though the long axis of the images they generate are along the East-West axis. It is therefore necessary to rotate the two sets of images by (almost) the same angle so that the direction of parallax matches the horizontal axis of the screen on which they will be projected.

Figure 4 illustrates the impact of the images alignment and rotation.

4 3D viewing

4.1 Stereoscopic techniques

The 3D images generated above have been tested with different stereoscopic techniques, which are listed below in the order of increasing hardware requirement and quality of the rendition.

⁹ www.mayaskies.net/productiontools/articles/Introducingthedome-master.pdf.




Fig. 4. Anaglyph rendition of an Aurora snapshot, before (upper panel) and after (lower panel) alignment of the images produced by the left and right cameras.

4.1.1 Free-viewing and cross-eye

The left and right images are shown side by side, either at their respective location and watching far beyond the screen, in the so-called free-viewing technique, as shown in Figure 5, or after swapping their position (i.e., left image on the right side and vice versa), and crossing the eyes, in the aptly named cross-eye technique. These techniques, requiring no material, are by far the cheapest and most accessible ones, but take some time to master, and are limited to very small images.

4.1.2 Anaglyphs

Since the Auroras are mostly green in color, a single composite image of the realigned observations is made in which the red channel is the green component of the left image, and the blue channel is made out of the right image. This is looked at with anaglyph glasses in which the left glass is red, and the right one is blue. This technique, available on screen and in print, as illustrated in Figure 4, is extremely cheap, with the drawback of losing color information and reducing luminosity.

4.1.3 Passively polarized screen and glasses

On an ad hoc screen, the odd lines, polarized vertically, show every second line of the left image while the even lines, polarized horizontally, do the same for the right image. The left and right lenses of a pair of glasses are polarized vertically and horizontally, respectively. More recent models of screens and glasses use circular polarization instead. This technique, used

on some 3D TV screens, is affordable and preserves the colors, but has a resolution cut in half.

4.1.4 Time multiplexing and active glasses

The left and right images are shown alternatively on a LCD screen polarized at 45° and supporting high frequency refreshment rate, and the left and right glasses become alternatively opaque and transparent to that polarization, thanks to rotating liquid crystals, with a shutter system operating in synch with the screen. This technique, preferred for computer video games, requires specific glasses and proprietary software and video card. Like all techniques based on linear polarization, it is limited to small or medium size screens.

4.1.5 Other techniques

Other techniques that we were not able to test with our Aurorae images, but are extremely suitable for very large screens, include time multiplexing of circularly polarized images, shown in alternance, with circularly polarized passive glasses, very widespread in 3D movie theaters, or wavelength multiplexing, where the left and right images use different red, green and blue wavelength bands, to which the left and right dichroic lenses are, respectively, transparent, and which is mostly used in planetariums.

Yet another pathway is the use of virtual reality glasses, allowing the immersion of the viewer in a outdoor scene illuminated by 3D Northern Lights. However, for maximum efficiency this requires a heavier observational set-up providing two different views of, at least, the whole hemisphericsky. Set-up that we are only starting to implement.

4.2 Depth of images

In all the techniques listed above, the stars, which after the image processing of Section 3 match exactly on the left and right images, will appear to be on the plan of the paper or of the screen, with the Aurora floating in front of them. Depending on the technique used and the size of the image, and in order to improve the feeling of immersion, it may be necessary to move the whole 3D image closer to or further from the viewer. This can be achieved by shifting, for instance, the left image toward the right or the left, respectively.

5 Conclusion

In this paper we present the image processing pipeline implemented to produce real time 3D movies of Aurorae Borealis out of images produced by two distant high-resolution fish-eye cameras. A model of the camera optical response is proposed, and for each camera, its eight numerical parameters are constrained using the positional astrometric information of the two dozen bright stars it observed. Each frame of the filmed rushes can be then corrected from this response in order to properly superpose the images, and produce 3D movies with a good stereoscopy. Samples of minutes-long movies produced by this pipeline are available for download at <http://www.iap.fr/aurora3d>, and can be watched on monitors or projection screens




Fig. 5. Left and right images of an Aurora, cropped to a 4:3 aspect ratio, for 3D free-viewing, as described in Section 4. The cross-eye technique, where the left and right images are swapped, can be tried by turning the image up-side down.

adapted for 3D vision, as described in Section 4, and probably in virtual reality headsets. The IDL source code of the pipeline, and its fish-eye projection model, will be made available at the same location, after it has been properly cleaned-up and documented.

We are now applying the same techniques to images and films produced with two pairs of camera, separated by ~ 6 km as previously, and the cameras in each pair set-up so that the whole hemispheric sky is observed at once by each pair. After careful stitching of the images, this provides higher quality 3D movies suitable for hemispherical screens, such as planetariums, or for virtual reality glasses, at the cost of heavier observational setup and logistics.

We hope that the existing movies, and the forthcoming ones, will be of interest to the general public raptured by the spectacle of Northern Lights, as well as to scientists studying the phenomenon.

Acknowledgments. We thank François Sèvre and Stéphane Colombi for insightful discussions, François Bouchet for his funding of this article publication charges, Inge Petter, Amlien and Elisabeth Nygård for their hospitality and friendship in the far North, as well as Urpo and Tarja Taskinen for their cozy welcome on the Arctic polar circle. The authors and the editor thank two anonymous referees for their very constructive comments and suggestions, which helped improving this paper and for their assistance in evaluating it.

Appendix A Impact of atmospheric refraction

The atmospheric refraction makes astronomical objects appear higher above the horizon (i.e., at smaller local colatitude, measured from the zenith) than they actually are, with an offset increasing with the colatitude. In the simplest model, in which the atmosphere is approximated as a uniform and flat slab of gas, a source of actual colatitude δ will appear at a colatitude δ_a shifted by

$$R_{\text{atm.}} = \delta - \delta_a = R_{45} \tan(\delta_a) \approx R_{45} \tan(\delta), \quad (\text{A.1})$$

and Comstock (1890) proposed that the amplitude of refraction, parameterized by its value at 45° , R_{45} , expressed in arcsec, depends on the observational conditions via

$$R_{45} = \frac{21.5 P}{273 + T}, \quad (\text{A.2})$$

with P and T the atmospheric pressure and temperature, in mmHg and $^\circ\text{C}$, respectively. One then notes that, for small values of $R_{\text{atm.}}$, the apparent equisolid radial distance is

$$r_{ES}(\delta - R_{\text{atm.}}) \simeq r_{ES} - R_{\text{atm.}} \cos(\delta/2), \quad (\text{A.3a})$$

$$= r_{ES} - R_{45} \frac{r_{ES}}{1 - \frac{r_{ES}^2}{4 - r_{ES}^2}}, \quad (\text{A.3b})$$

$$= (1 - R_{45}) r_{ES} - R_{45} \sum_{p \geq 1} \frac{r_{ES}^{2p+1}}{2^{p+1}}. \quad (\text{A.3c})$$

At sea level, for high pressure ($P = 800$ mmHg = 107 kPa) and low temperature ($T = -20^\circ\text{C}$), $R_{45} = 68'' = 3.3 \times 10^{-4}$ rad, and the apparent location of a source seen through a fisheye of focal f with perfect equisolid projection would be

$$r(\delta)/f = 0.99967 r_{ES} - 8.24 \cdot 10^{-5} r_{ES}^3 - 4.12 \cdot 10^{-5} r_{ES}^5 + \dots \quad (\text{A.4})$$

Of course, equation (A.1) is very crude model of atmospheric refraction, but for sources 10° or more above the horizon, it only differs by a few arcsec at most from more accurate and complex models. Moreover, this model most likely over-estimates the refraction close to the horizon since it diverges at $\delta = 90^\circ$, leading us to consider that the actual refraction will be smaller than the one estimated above.

Appendix B Discretization error

Since the projection of the optics is close to equisolid, as illustrated in Figure 2, it preserves the areas and therefore the

surface densities. If the distribution of stars is assumed uniform on the sky, it will also be uniform on the image. When assigning a star position (x, y) to the closest pixel of integer indices (i, j) the error made is $(dx, dy) = (x - i, y - j)$, where dx and dy both have a probability density of 1 on the interval $] -1/2, 1/2[$. Therefore, the means and variances of these errors are $\langle dx \rangle = \langle dy \rangle = 0$ and $\langle (dx)^2 \rangle = \langle (dy)^2 \rangle = 1/12$. Defining the actual and approximate radial distances as $r \equiv (x^2 + y^2)^{1/2}$ and $r' \equiv (i^2 + j^2)^{1/2}$, respectively, then their difference

$$dr = r - r', \quad (\text{B.1a})$$

$$\simeq \frac{(dx)^2 + (dy)^2}{2r'} + \frac{idx + jdy}{r'} \quad \text{for } r' \gg 1, \quad (\text{B.1b})$$

has for mean value $\langle dx \rangle = 1/(12r')$ and for standard deviation

$$\sigma_{dr} = 1/\sqrt{12} \simeq 0.28867. \quad (\text{B.2})$$

One also notes that $s \equiv (dx)^2$ has the density for $0 \leq s \leq 1/4$, and therefore

$$e^2 \equiv (dx)^2 + (dy)^2, \quad (\text{B.3})$$

has the density

$$\rho_{e^2} = \pi \quad \text{for } 0 \leq e^2 \leq 1/4, \quad (\text{B.4a})$$

$$= 2\arcsin\left(\frac{1}{2e^2} - 1\right) \quad \text{for } 1/4 < e^2 \leq 1/2, \quad (\text{B.4b})$$

$$= 0 \quad \text{otherwise.} \quad (\text{B.4c})$$

It implies

$$\langle e \rangle = \left(\sqrt{2} + \ln(1 + \sqrt{2})\right)/6 \simeq 0.3826, \quad (\text{B.5})$$

and

$$\sigma = (1/6 - \langle e \rangle^2)^{1/2} \simeq 0.14242. \quad (\text{B.6})$$

References

- Borovička J, Spurný P, Keclíková J. 1995. A new positional astrometric method for all-sky cameras. *Astron Astrophys Suppl* **112**: 173–178.
- Comstock C. 1890. A simple approximate formula for the refraction. *Sidereal Messenger* **9**: 185–190.
- Coupinot G, Hecquet J, Auriere M, Futaully R. 1992. Photometric analysis of astronomical images by the wavelet transform. *Astron Astrophys* **259**: 701–710.
- Hallinan TJ, Stenbaek Nielsen HC, Deehr CS. 1985. Enhanced aurora. *J Geophys Res: Space Phys* **90**(A9): 8461–8475.
- Kataoka R, Fukuda Y, Uchida HA, Yamada H, Miyoshi Y, Ebihara Y, Dahlgren H, Hampton D. 2016. High-speed stereoscopy of aurora. *Annal Geophys* **34**(1): 41–44.
- Kataoka R, Miyoshi Y, Shigematsu K, Hampton D, Mori Y, Kubo T, Yamashita A, Tanaka M, Takahei T, Nakai T, Miyahara H, Shiokawa K. 2013. Stereoscopic determination of all-sky altitude map of aurora using two ground-based Nikon DSLR cameras. *Annal Geophys* **31**: 1543–1548. doi: [10.5194/angeo-31-1543-2013](https://doi.org/10.5194/angeo-31-1543-2013).
- Markwardt CB. 2009. Non-linear least-squares fitting in IDL with MPFIT. In: Bohlender DA, Durand D, Dowler P, eds. *Astronomical data analysis software and systems XVIII*, Astronomical Society of the Pacific Conference Series, Vol. 411, p. 251.
- Marquardt DW. 1963. An algorithm for least-squares estimation of nonlinear parameters. *J Soc Indust Appl Math* **11**(2): 431–441.
- Marr D, Hildreth E. 1980. Theory of edge detection. *Proc R Soc Lond B: Biol Sci* **207**(1167): 187–217.
- Meeus J. 1991. *Astronomical algorithms*. Willmann-Bell, Richmond.
- More JJ. 1978. The Levenberg-Marquardt algorithm: implementation and theory. In: Watson GA ed, *Numerical Analysis, Proceedings of the Biennial Conference Held at Dundee, June 28–July 1, 1977*. Springer, Berlin Heidelberg, pp. 105–116. doi: [10.1007/BFb0067700](https://doi.org/10.1007/BFb0067700).
- Mori Y, Yamashita A, Tanaka M, Kataoka R, Miyoshi T, Okutomi M, Asama H. 2013. Calibration of fish-eye stereo camera for aurora observation. In: *Proceedings of the International Workshop on Advanced Image Technology (IWAIT2013)*, January 2013, pp. 729–734.
- Press W, Teukolsky S, Vetterling W, Flannery B. 1992. *Numerical recipes in Fortran: the art of scientific computing*, 2nd ed. Cambridge University Press, New York. ISBN: 052143064X.
- Stephenson FR, Willis DM, Hallinan TJ. 2004. The earliest datable observation of the aurora borealis. *Astron Geophys* **45**(6): 6.15–6.17. doi: [10.1046/j.1468-4004.2003.45615.x](https://doi.org/10.1046/j.1468-4004.2003.45615.x).