Optical Flow Computation in the Log-Polar Plane

Konstantinos Daniilidis and Volker Krüger*

Computer Science Institute, Christian-Albrechts University Kiel
Preusserstr. 1-9, 24105 Kiel, Germany
email: kd@informatik.uni-kiel.de

Abstract. The goal of this paper is to study how the logarithmic-polar mapping affects the spatiotemporal volume and the computation of the optical flow. We propose two new methods for the estimation of the optical flow and its spatial derivatives in the log-polar plane. We study analytically and experimentally the effects of the polar deformation and we qualitatively explain the decimation due to subsampling on the computation of optical flow.

1 Introduction

This paper is concerned with the computation of optical flow in image sequences obtained with the logarithmic polar transformation. The log-polar transformation is a model for space-variant resolution in the periphery of the image. Space-variant sensing arises as a necessity in systems which must be able to process simultaneously a central region of interest (fovea) in detail for recognition tasks and a wide-angle peripheral view for detecting events and new candidates for gaze change. Uniform resolution in the peripheral part would result to a computational burden unacceptable for real-time reactive behavior. The biologically motivated log-polar transformation has a second significant advantage: It is a very rich representation regarding recognition tasks (rotation and scaling invariance) as well as navigational tasks (ego-motion and time to collision estimation, motion detection).

The goal of this paper is to study what kind of information which can be used for motion tasks is still preserved after the log-polar transformation . We study analytically and experimentally the effects of the polar deformation in sec. 3 and we qualitatively explain the logarithmic subsampling in sec. 4. In particular,

- we prove that the polar transformation introduces fictitious gray-values curvature that leads to an erroneous elimination of the aperture problem,
- we propose two new methods for the optical flow estimation in the log-polar domain that are superior to methods directly transferred from the cartesian domain and we experimentally study their performance in a real sequence,
- we propose a basis for an analysis of the logarithmic subsampling that allows spectral techniques for the design of the necessary low-pass and gradient filters.

* On leave at Lira-Lab, DIST, University of Genova, Via all'Opera Pia 13, 16145 Genova. We greatfully thank Gerald Sommer, Giulio Sandini, Yehoshua Zeevi, and Udo Mahlmeister for the helpful discussions.

We use (x, y) for the cartesian coordinates and (ρ, η) for the polar coordinates in the plane. We apply the log-polar mapping on the non-foveal part of a retinal image. Therefore, we define as the domain of the mapping the ring-shaped area $\rho_0 < \rho < \rho_{max}$ where ρ_0 and ρ_{max} are the radius of the fovea and the half-size of the retinal image, respectively. Furthermore, a hardware CCD-sensor with the log-polar property or a software implementation of the mapping needs a discretization of the log-polar plane. By assuming that N_r is the number of cells in the radial direction and N_a is the number of cells in the angular direction the mapping from the polar coordinates (ρ, η) to the log-polar coordinates (ξ, γ) reads (see also [8]) $\xi = \log_a(\frac{\rho}{\rho_0}), \quad \gamma = \frac{N_a}{2\pi}\eta$, where the logarithmic basis a is obtained from the foveal radius ρ_0, the image radius ρ_{max} and the radial resolution N_r: $a = \exp(\frac{1}{N_r} ln(\frac{\rho_{max}}{\rho_0}))$. From now on we will use only η ranging from 0 to 2π. The mapping of the gray-value function $I(x, y)$ in the cartesian plane to the gray-value function $J(\xi, \eta)$ in the log-plane is by no means trivial. Every log-polar cell corresponds to a receptive field in the cartesian plane[2]. The image $J(\xi, \eta)$ is the result of a space-variant filtering that affects all subsequent computations on the log-polar plane like spatiotemporal filtering appearing later in this paper. We will not delve in this issue here. It has been extensively studied in [1] but it still remains an open problem as we will see in section 4. In our implementation we used non-overlapping averaging receptive fields as implemented in the emulation of the space-variant sensor in [8].

2 Optical flow in the log-polar plane

We will use the notion of optical flow for the apparent velocity of gray-value structures in the image as opposed to the pure geometric definition of the motion field as the velocity of the projected scene points on the image. Interpretations and actions concerning the scene are based on the motion field although only the optical flow field can be observed. This discrimination becomes more crucial here than in the cartesian plane due to the polar deformation and the logarithmic subsampling of the gray-value function. Using the coordinate transformation we are able to transform exactly the cartesian motion field on the log-polar plane. However, the deformation of the gray-value function causes new apparent shifts of the gray-value function or eliminates existing ones.

In this section, we first apply flow computation methods already existing for the cartesian plane to a log-polar image sequence. We denote by (u, v) the optical flow vector in the cartesian plane and by (u^l, v^l) the optical flow vector in the log-polar plane. The motion field vectors in the cartesian and log-polar plane are denoted by $(\dot{x}, \dot{y})$ and $(\dot{\xi}, \dot{\eta})$, respectively. We first compute the motion field vectors in the polar plane. The definition of the polar and log-polar coordinates yields

$$\dot{\rho} = \dot{x}\cos\eta + \dot{y}\sin\eta \qquad \dot{\xi} = \frac{1}{\ln a}\frac{\dot{\rho}}{\rho} \qquad \dot{\eta} = \frac{1}{\rho}(-\dot{x}\sin\eta + \dot{y}\cos\eta). \qquad (1)$$

[2] There is a large amount of results on biological findings the space of this article does not allow to cite [7; 2]

To compute optical flow we will use methods based on the spatiotemporal derivatives of the image and on the Brightness Change Constraint Equation (BCCE) on the log-polar image $J(\xi, \eta)$:

$$J_\xi u^l + J_\eta v^l + J_t = 0, \tag{2}$$

where J_ξ, J_η, and J_t are the spatiotemporal derivatives of the image. A first method is the application of the BCCE in the neighborhood $(\xi + \delta\xi, \eta + \delta\eta)$ of every considered pixel (ξ, η) assuming that the optical flow is locally constant [3]. We give appropriate weights $w(\delta\xi, \delta\eta)$ to the application of BCCE at every pixel so that the influence is higher in the center of the neighborhood. The solution is obtained by minimizing

$$\sum_{\delta\xi, \delta\eta} w(\delta\xi, \delta\eta)^2 \left(J_\xi(\xi+\delta\xi, \eta+\delta\eta) u^l + J_\eta(\xi+\delta\xi, \eta+\delta\eta) v^l + J_t(\xi+\delta\xi, \eta+\delta\eta) \right)^2 \tag{3}$$

with respect to (u^l, v^l). The above minimization problem is equivalent to the weighted minimization of the form $\|W(A\boldsymbol{u^l} - b)\|$ which is solved by singular value decomposition.

The second method is an extension of the first one. It allows the linear variation of the flow inside the neighborhood enabling, thus, an estimation of flow as well as its spatial derivatives [6]. The BCCE equation with linearly varying flow is applied for every $(\delta\xi, \delta\eta)$ inside a neighborhood yielding an overconstrained system with six unknowns. However, both these already known methods use assumptions about local constancy or affinity of the optical flow that do not reflect the harmonic variation of both flow components (1) with the angle η in the log-polar plane. We will delve into this problem in the next section. We finish this section giving abbreviations to the presented methods. As of now we will call LCT the method based on the Local Constancy of the flow in the Transformed image (polar or log-polar) and LAT to the method based on the Local Affinity of the flow in the Transformed image.

3 The polar deformation

The first source of error in the log-polar optical flow field is due to the polar deformation of the gray-value function. The polar transformation maps straight edges into curved edges (Fig. 1) enabling thus the computation of both components of the optical flow at points without curvature in the original cartesian image. This superficial elimination of the aperture problem introduces optical flow values with a large error regarding the expected motion field. In this section we will first transform the neighborhood–gradient approach into a second derivative method in order to study analytically the rank of the resulting linear system. Then we will prove the expected fact that the linear system in the polar plane has full rank even if the Hessian matrix in the cartesian plane is singular. The aperture problem is always eliminated by introducing some assumption on the local variation of the flow. In this sense, we assume the local constancy of the back-transformed flow in the cartesian plane. We, thus, treat the aperture

problem in the cartesian plane. The resulting matrix of gray-value derivatives in the polar plane will be proved to have the same rank as the cartesian Hessian.

We denote by $E(\rho,\eta)$ the gray-value function and by (u^p, v^p) the optical flow in the polar domain. Using the Brightness Change Constraint Equation and the assumption that the polar flow is locally constant and applying the Taylor expansion to the derivatives at the positions $(\rho+\delta\rho, \eta+\delta\eta)$ yields the overconstrained system

$$E_\rho u^p + E_\eta v^p + E_t = 0 \quad E_{\rho\rho} u^p + E_{\eta\rho} v^p + E_{t\rho} = 0 \quad E_{\rho\eta} u^p + E_{\eta\eta} v^p + E_{t\eta} = 0 \quad (4)$$

where we have omitted the resulting weights on each equation. We are interested in the 2×2 coefficient matrix of the second and third equation which is the Hessian of the polar gray-value function $E(\rho,\eta)$. By twice differentiating $E(\rho,\eta)$ it can be easily proved that

$$\begin{pmatrix} E_{\rho\rho} & E_{\eta\rho} \\ E_{\rho\eta} & E_{\eta\eta} \end{pmatrix} = \frac{\partial(x,y)}{\partial(\rho,\eta)}^T \begin{pmatrix} I_{xx} & I_{xy} \\ I_{xy} & I_{yy} \end{pmatrix} \frac{\partial(x,y)}{\partial(\rho,\eta)} + I_x \begin{pmatrix} \frac{\partial^2 x}{\partial\rho^2} & \frac{\partial^2 x}{\partial\rho\partial\eta} \\ \frac{\partial^2 x}{\partial\rho\partial\eta} & \frac{\partial^2 x}{\partial\eta^2} \end{pmatrix} + I_y \begin{pmatrix} \frac{\partial^2 y}{\partial\rho^2} & \frac{\partial^2 y}{\partial\rho\partial\eta} \\ \frac{\partial^2 y}{\partial\rho\partial\eta} & \frac{\partial^2 y}{\partial\eta^2} \end{pmatrix} \quad (5)$$

All the matrices with derivatives of (x,y) with respect to (ρ,η) have full rank up to two values of η for the second derivatives. Hence, the singularity of the cartesian Hessian does not lead to the singularity of the polar Hessian which also depends on the cartesian gradient. It is plausible to suppose that the smallest singular value of the coefficient matrix – used as a confidence measure – will be higher in the system (4) than in the equivalent system in the cartesian plane. This fact causes the acceptance of erroneous optical flow values.

We proceed by substituting the assumption of local flow constancy in the polar plane with local flow constancy before applying the polar transformation. We differentiate the Brightness Change Constraint Equation assuming that the spatial derivatives of (u,v) vanish and we obtain a system with new coefficient matrix

$$\begin{pmatrix} E_{\rho\rho} & E_{\eta\rho} - \frac{E_\eta}{\rho} \\ E_{\rho\eta} - \frac{E_\eta}{\rho} & E_{\eta\eta} + \rho E_\rho \end{pmatrix} = \frac{\partial(x,y)}{\partial(\rho,\eta)}^T \begin{pmatrix} I_{xx} & I_{xy} \\ I_{xy} & I_{yy} \end{pmatrix} \frac{\partial(x,y)}{\partial(\rho,\eta)} \quad (6)$$

which proves that the singularity of the cartesian Hessian matrix is the necessary and sufficient condition for the singularity of the new coefficient matrix. Hence, the coefficient matrix (6) does not introduce erroneous values of the optical flow like the matrix (5) in the system (4).

Based on this fact we are going to construct a gradient-sampling method like (3). We transfer the assumptions about constancy and affinity of the *cartesian* flow to the application of the BCCE in the neighborhood pixels $(\xi+\delta\xi, \eta+\delta\eta)$ allowing the log-polar flow to vary

$$\begin{pmatrix} u^l \\ v^l \end{pmatrix} = \frac{1}{\rho_0 a^{\xi+\delta\xi}} \begin{pmatrix} \frac{\cos(\eta+\delta\eta)}{\ln a} & \frac{\sin(\eta+\delta\eta)}{\ln a} \\ -\sin(\eta+\delta\eta) & \cos(\eta+\delta\eta) \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}. \quad (7)$$

We will call this method the Local Constancy in the Cartesian image (LCC) method. The next step is straightforward. We allow the cartesian flow to vary

linearly in the local neighborhood. Combining Cartesian local linear variation with the Log-polar BCCE (2), we get

$$J_t = -\frac{1}{\rho a^{\xi+\delta\xi}} \begin{pmatrix} J_\xi & J_\eta \end{pmatrix} \begin{pmatrix} \frac{\cos(\eta+\delta\eta)}{\ln a} & \frac{\sin(\eta+\delta\eta)}{\ln a} \\ -\sin(\eta+\delta\eta) & \cos(\eta+\delta\eta) \end{pmatrix} \begin{pmatrix} u + u_x\delta x + u_y\delta y \\ v + v_x\delta x + v_y\delta y \end{pmatrix} \quad (8)$$

We call this method the Local Affinity in the Cartesian image (LAC) method. By setting $\xi = 0$ and $\eta = 0$ we obtain the equations derived in [9] as a special case.

4 The logarithmic subsampling

In this section we present first steps towards understanding the effects of the logarithmic subsampling on the computation of the spatiotemporal derivatives. We restrict our study on a 1D gray-value function $g(\rho)$ on the discrete domain $\rho = \rho_0..\rho_{max}$ and its logarithmically subsampled version $\gamma(\xi)$ defined on $\xi = 0..N$ as if both were the radius and its log-polar map of a 2D-image without polar subsampling. We introduce an intermediate function $\lambda(\xi)$ [5] obtained by exact coordinate transformation $\lambda(\xi) = g(\rho_0 b^\xi)$ where the logarithm's basis b is chosen in such a way that the discrete original signal is transformed without loss. This means that the basis b must be less equal than the coordinates ratio $b=\rho_{max}/(\rho_{max}-1)$ so that even the gray-value of the most peripheral pixel is exactly transformed. The lossless signal is sparse and has a dimension $M = \ln(\rho_{max}/\rho_0)/\ln b$ much greater than the original signal so we interpolate for the intermediate valueless pixels. The logarithmic signal $\gamma(\xi)$ is then obtained by the three steps of linear shift-invariant low pass filtering, subsampling, and shrinking as in a layer transition step in a regular pyramid. The subsampling interval is M/N where M is the dimension of the lossless signal as above and $N = \ln(\rho_{max}/\rho_0)/\ln a$ is the final resolution of the logarithmic signal $\gamma(\xi)$.

The introduction of the lossless image enables the study of the decimation effects with spectral methods. First, we become able to design the appropriate low-pass filters to suppress the energy above half of the subsampling frequence (future work). Second, considering the optical flow u_λ in the lossless signal as the flow with the lowest error we can use it as the reference for the error introduced by subsampling in the flow u_γ. If we use a differential technique as in the previous sections the flow of the 1D lossless and subsampled signals is $u_\lambda = -\frac{\lambda_t}{\lambda_\xi}$ and $u_\gamma = -\frac{\gamma_t}{\gamma_\xi}$, respectively.

The temporal gradient γ_t of the logarithmic signal is the response of the lossless temporal gradient to the low-pass filtering step. To study the spatial gradient we first point out that the gradient is a bandpass e.g. the first derivative of a Gaussian here. We assume that the low pass filter has the appropriate antialiasing characteristics. The spectrum of the lossless image is stretched out by the subsampling and shrinking steps whereas the spectrum of the spatial gradient is the same before and after subsampling. The error in the spatial gradient depends on the amplification or attenuation of frequency contributions according to the contributions under the frequency support of the gradient. We

conjecture that such a decreasing of γ_ξ might be the reason for the observed systematic overestimation of the length of the optical flow γ_t/γ_ξ.

5 Experimental comparison

In this section we present results that compare the four flow computation methods described and show the effect of the polar deformation before the logarithmic subsampling. In all figures the log-polar images are drawn such that the η-axis is the horizontal axis and the ξ-axis is the vertical axis pointing downwards. To interpret the log-polar images we note that the angle η is measured beginning counterclockwise from the y-axis that is pointing downwards. So moving horizontally in the log-polar plane we first see the transformed lower right quadrant, then the transformed upper right quadrant and so on. The compression rates obtained by the log-polar transformation are about 1:25 for all analyzed sequences. The spatiotemporal derivatives are computed by convolving with the binomial differentiation filters. The local constancy and the local affinity assumption are applied in 5×5 neighborhoods. The least squares problems are solved with the Singular Value Decomposition and a threshold is applied to the smallest singular value as a reliability criterion.

We first tested the Local Constancy methods on the polar transform of an image sequence consisting of black squares on white background moving with uniform velocity of (1,1). The polar transformed image is shown on the left of Fig. 1. The images on the middle and left of Fig. 1 show the smallest singular value of the linear systems using the Local Constancy assumption in the Polar (LCT) and the Cartesian (LCC) image, respectively. As we already expected from eq. 7 the LCC method produces a system with lower singular values.

To test the methods we used the real sequence "Marbled Block" [3] [4] with known ground-truth values for the motion field. The original image of the sequence and its log-polar transform are shown in Fig. 2 left and middle, respectively. In Fig. 2 (right) we show the flow field computed with the LAC method.

The error measures used are the relative error and the angle between $(u, v, 1)$ and $(\hat{u}, \hat{v}, 1)$ where (u, v) and $(\hat{u}, \hat{v})$ are the ground-truth and estimated flow, respectively. Furthermore, we compare not only to the transformed ground-truth flow (left column in the tables) but also to the transformation of a flow estimated first in the cartesian plane (right column). The error with respect to the latter should be considered as a lower bound for the error expected. The density is the fraction of the estimates with smallest singular value above a threshold which varies for the four estimation techniques.

We first present (Tab.1) the angle- and relative errors for the *polar* transform of the sequence obtained with angular resolution of 512 samples per 360 degrees and radial resolution equal to the original (256). Regarding the local constancy assumption applied on the polar (LCT) and the cartesian (LCC) plane the errors are about the same for the same density. However, this density is achieved for appropriately chosen high threshold for the LCT method. The superiority of the LCC method is shown if we compare it to the performance of the LCT method

[3] Created by Michael Otte at University of Karlsruhe and FhG-IITB, Germany.

with the same threshold (LCT-thr).

	Transformed ground truth			**Transform of the cartesian estimate**		
Technique	av. ang. err.	av. rel. err.	density	av. ang. err.	av. rel. er	density
LCT	5.64729	18.74557	0.54977	3.68909	12.36976	0.54977
LCT-th	7.26917	23.65154	0.77566	3.78853	13.56477	0.77566
LCC	6.02655	19.11052	0.48914	3.63529	12.04593	0.48914

Table 1. Error statistics for the polar transform of the "Marbled Block" sequence (see text for explanation).

The log-polar transform of the "Marbled Block" sequence is obtained with angular resolution of 128 samples/360 degrees and radial resolution of 45 samples for the radial range [32..256]. It should be noted that the angular resolution is one-fourth of the angular resolution of the polar transform, therefore the error is due to both the logarithmic and the polar subsampling. The errors with respect to the transformed ground-truth values are not significantly higher than to the transformed estimates what means that the effects of noise and poor gray-value structure in the original are inferior to the subsampling effects. The errors are shown for about the same density in order to compare the performance at points where the coefficient matrices are regular in all methods. We are thus strict to the methods using assumptions in the cartesian plane. Even with such a high relative error we can use the log-polar transform for 3D-analysis: We show in a companion paper in these proceedings that the 3D-translation direction can be computed with only 5 degrees error.

	Transformed ground truth			**Transform of the cartesian estimate**		
Technique	av. ang. err.	av. rel. er	density	av. ang. err.	av. rel. er	density
LCT	5.34357	34.53929	0.71908	4.26153	32.02543	0.71908
LCC	5.79178	38.03186	0.56144	4.80384	34.88554	0.56144
LAT	4.33591	28.51493	0.79297	3.32836	26.21737	0.79297
LAC	4.30630	26.69212	0.72917	3.29727	23.21143	0.72917

Table 2. Error statistics for the log-polar transform of the "Marbled Block" sequence (see text for explanation).

References

1. M. Bolduc and M.D. Levine. A foveated retina system for robotic vision. In *ECCV-94 Workshop on Natural and Artificial Visual Sensors*, 1994.
2. R.E. Kronauer and Y.Y. Zeevi. Reorganization and Diversification of Signals in Vision. *IEEE Trans. Systems, Man, and Cybernetics*, 15:91–101, 1985.
3. B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *DARPA Image Understandig Workshop*, pp. 121–130, 1981.
4. M. Otte and H.-H. Nagel. Optical flow estimation: advances and comparisons. In *Proc. Third European Conference on Computer Vision*, pp. 51–60. Stockholm, Sweden, May 2-6, 1994.

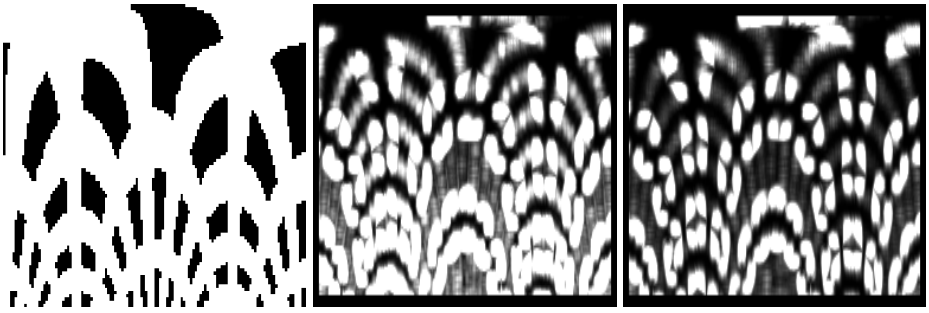

Fig. 1. The polar transform of an image with black rectangles (left), and the smallest singular value of the linear system resulting from the LCT (middle) and the LCC (right) assumptions, respectively. The horizontal axis is the angular axis and the vertical axis is the radial axis.

Fig. 2. The cartesian original image 512x512 pixels(left), its log-polar transform 77x128 pixels shown magnified (middle), and the computed optical flow (right) of the "Marbled Block" sequence. The horizontal axis is the angular axis and the vertical axis is the logarithmic radial axis.

5. M. Porat and Y.Y. Zeevi. The Generalized Gabor Scheme of Image Representation in Biological and Machine Vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10:452–468, 1988.
6. R. Schalkoff and E. McVey. A model and tracking algorithm for a class of video targets. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 4:2–10, 1982.
7. E.L. Schwartz. Spatial mapping in the primate sensory projection: analytic structure and relevance to perception. *Biological Cybernetics*, 25:181–194, 1977.
8. M. Tistarelli and G. Sandini. Dynamic aspects in active vision. *CVGIP: Image Understanding*, 56:108–129, 1992.
9. H. Tunley and D. Young. First order optical flow from log-polar sampled images. In *Proc. Third European Conference on Computer Vision*, pp. 132–137. Stockholm, Sweden, May 2-6, 1994.

This article was processed using the LaTeX macro package with LLNCS style