

Quasi-Periodic Spatiotemporal Filtering

Gertjan J. Burghouts and Jan-Mark Geusebroek

Abstract—This paper presents the online estimation of temporal frequency to simultaneously detect and identify the quasiperiodic motion of an object. We introduce color to increase discriminative power of a reoccurring object and to provide robustness to appearance changes due to illumination changes. Spatial contextual information is incorporated by considering the object motion at different scales. We combined spatio-spectral Gaussian filters and a temporal reparameterized Gabor filter to construct the online temporal frequency filter. We demonstrate the online filter to respond faster and decay faster than offline Gabor filters. Further, we show the online filter to be more selective to the tuned frequency than Gabor filters. We contribute to temporal frequency analysis in that we both identify (“what”) and detect (“when”) the frequency. In color video, we demonstrate the filter to detect and identify the periodicity of natural motion. The velocity of moving gratings is determined in a real world example. We consider periodic and quasiperiodic motion of both stationary and nonstationary objects.

Index Terms—Color, time-frequency analysis, video signal processing.

I. INTRODUCTION

THE temporal frequency of a moving object may be an important property of that object. Real world applications illustrate this, for instance when monitoring the oscillatory beating of a heart. Further, for periodically moving objects, the temporal frequency of the periodic motion directly relates to the velocity of the motion [1]. The velocity of waves propagating through water follows directly from its motion periodicity and its spatial frequency [2]. The velocity of waves is a direct consequence of an harmonic mechanical system, described by the wind force and the depth, width, and mass of the water, which is in equilibrium. The measurement of an object’s periodic motion hence may enable the estimation of both the object’s velocity and environmental properties derived thereof. Estimating velocity from motion periodicity is robust, since periodicity is invariant to the object’s distance. On the contrary, estimated motion from optical flow [3] varies with the object’s distance. In addition, periodic motion has proven to be an attentional attribute [4], which may facilitate target detection in video (see, e.g., visual surveillance in [5]).

To measure the periodicity of object motion, we propose a temporal frequency filter that measures the reoccurrence of an object’s surface during a time interval. Note that the class of periodic temporal events is more rigid than the class of stochas-

tically defined dynamic textures [6]. The temporal frequency filter cannot measure both the frequency and the timing of an occurrence of periodic motion with arbitrary precision [7]. The challenge for detecting and identifying temporal frequency is thus to find the right trade-off between timing and frequency analysis. Time-frequency analysis based on the Fourier transform of the video signal [5], [8], [9] ignores temporal discrimination. However, the Fourier transform extracts maximum information about the frequency composition of the signal. Gabor filtering provides the optimal joint resolution in both time and frequency, obtaining equal temporal width at all frequencies [7], [10]. Hence, the Gabor temporal frequency filter measures both the frequency identification (“what”) and the frequency detection (“when”).

We embed a temporal frequency filter in the Gaussian scale-space paradigm [11] to incorporate the spatial and temporal scale in its measurement. Larger spatial scales incorporate contextual information, hence avoiding pixel matching. A temporal scale allows the periodicity of object motion to be resolved in suitable time windows. For the analysis of temporal frequency, it is natural to measure the temporal signal in the Fourier domain. A Gaussian measurement in the Fourier domain, tuned to a particular frequency, boils down to a Gabor measurement [7] in the temporal domain. For online filtering, only the past is available. We deal with this restriction by a logarithmical mapping of the filter onto the “past” only [12]. However, the sinusoidal sensitivity curve of the temporal Gabor filter becomes logarithmical hence not suitable for frequency measurements. We reparameterize the temporal Gabor filter to optimize it for the local and online measurement of temporal frequency. We introduce color to increase discriminative power when measuring the reoccurrence of a particular surface.

In this paper, we derive an online temporal frequency filter and demonstrate the filter to respond faster and decay faster than Gabor filters. Additionally, we show the online filter to be more selective to the tuned frequency than Gabor filters (Section II). In color video, the filter detects and identifies the periodicity of natural motion. Further, we determine the velocity of moving gratings in a real world example (Section III). We demonstrate the general applicability of the proposed filter. Consequently, we do not attribute specialized topics that analyze motion of specific kinds in depth, such as motion-based recognition [1], [8], [9]. The experiments include a) stable and changing periodic motion of b) stationary and nonstationary objects with c) smooth and regularly textured surfaces.

II. TEMPORAL FREQUENCY FILTER

A. Derivation

We consider color video to be an energy distribution over space, wavelength spectrum and time. A spatio-spectral energy

Manuscript received November 23, 2004; revised May 11, 2005. This work was supported by the Netherlands Organization for Scientific Research (NWO). The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Philippe Salembier.

The authors are with the Intelligent Systems Lab of Amsterdam, Department of Computer Science, Faculty of Science, University of Amsterdam, 1098 SJ Amsterdam, The Netherlands (e-mail: burghout@science.uva.nl; mark@science.uva.nl).

Digital Object Identifier 10.1109/TIP.2005.864234

distribution is only measurable at a certain spatial resolution and a certain spectral bandwidth [13], [14]. Analogously, the temporal energy distribution is only measurable at a certain temporal resolution. Hence, physical realizable color video measurements inherently imply integration over spectral, spatial and temporal dimensions. Based on linear scale space assumptions [11], we consider Gaussian filters and their derivatives to measure color video. We generally define an i th order Gaussian derivative filter $G_{a^i}(a)$ probing a variable a at scale σ_a and location a_0

$$G_{a^i}^{\sigma_a, a_0}(a) = \frac{\partial^i G^{\sigma_a, a_0}(a)}{\partial^i a} = \frac{H_i\left(\frac{a-a_0}{\sigma_a}\right)}{\sqrt{2\pi}\sigma_a} e^{-\frac{(a-a_0)^2}{2\sigma_a^2}} \quad (1)$$

where the i th order Hermite polynomial with respect to $(a - a_0)/\sigma_a$, $H_i((a - a_0)/\sigma_a)$ determines the shape of the i th order Gaussian derivative filter. For orders $i \in \{1, 2, 3\}$, the Hermite polynomials H_i are given by $\{2(a - a_0)/\sigma_a, 4((a - a_0)/\sigma_a)^2 - 2, 8((a - a_0)/\sigma_a)^3 - 12(a - a_0)/\sigma_a\}$. For notational convenience, we omit the scale and location parameters where possible.

An object's surface is defined by its reflectance function $R(x, y, \lambda)$ at a spatial location (x, y) , where λ denotes the wavelength [13]. Furthermore, the temporal periodicity of the object is measured in time t [15]. The temporal frequency measurement hence requires a simultaneous measurement of these variables to determine whether an object's surface has re-occurred at a certain spatial location. The periodic reoccurrence of an object's surface at a constant time period p is defined as

$$\hat{E}(x, y, \lambda, t) = \hat{E}(x + x', y + y', \lambda, t + p) \quad (2)$$

with \hat{E} the measurement of the color video signal and (x', y') the translation of the point due to object movement relative to the camera. In the sequel, we consider the temporal frequency measurement at a spatial location (x, y) , and correct for the object's translation by tracking the object. The temporal frequency measurement $\hat{E}(x, y, \lambda, t)$ of the color video signal $E(x, y, \lambda, t)$ is performed by a filter $F(x, y, \lambda, t)$, yielding

$$\hat{E}(x, y, \lambda, t) \equiv E(x, y, \lambda, t) * F(x, y, \lambda, t) \quad (3)$$

with $(*)$ the convolution operator as we consider linear measurements.

For convenience, we first concentrate on the measurement of the wavelength distribution. To measure wavelength in color video, we consider the advantage to separate the luminance from the color channels. The opponent color system used in this paper is formalized by measuring with three spectral Gaussian derivative filters [14]: G_{λ^i} . The zeroth order derivative filter measures the energy over all wavelengths (the luminance), whereas the first order derivative filter compares the first half (blue) and second half (yellow) of the spectrum and the second order derivative filter compares the middle (green) and two outer (red) regions of the spectrum. To obtain colorimetry with human vision, the Gaussian filters are to be tuned such that the filters span the same spectral subspace as spanned by the CIE 1964 XYZ sensitivity curves. The location λ_0 and scale σ_λ of the Gaussian spectral filters are optimized such that

approximate colorimetry is obtained by setting the parameters to $\sigma_\lambda = 55$ nm, and $\lambda_0 = 540$ nm [14]

$$G_{\lambda^i}^{\sigma_\lambda=55 \text{ nm}, \lambda_0=540 \text{ nm}}(\lambda), i \in \{0, 1, 2\}. \quad (4)$$

See Fig. 1 for the sensitivity curves of the spectral filters. Spectral derivative filters $G(\lambda)$, $G_\lambda(\lambda)$, and $G_{\lambda\lambda}(\lambda)$ yield, respectively, the measurements \hat{E} , \hat{E}_λ , and $\hat{E}_{\lambda\lambda}$. In practice, the values are obtained by a linear combination of given RGB sensitivities [14]

$$\begin{bmatrix} \hat{E} \\ \hat{E}_\lambda \\ \hat{E}_{\lambda\lambda} \end{bmatrix} = \begin{pmatrix} 0.06 & 0.63 & 0.27 \\ 0.30 & 0.04 & -0.35 \\ 0.34 & -0.60 & 0.17 \end{pmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \quad (5)$$

Color can only be measured by integration over a spatial area and a spectral bandwidth. Hence, a color measurement requires a combination of the spectral filter (4) and a spatial filter. For simplicity, we select a zeroth order, isotropic two-dimensional (2-D) spatial filter [11]

$$G^{\sigma_{xy}}(x, y) = \frac{1}{2\pi\sigma_{xy}^2} e^{-\frac{x^2+y^2}{2\sigma_{xy}^2}} \quad (6)$$

where σ_{xy} indicates the spatial extent of the filter. To measure elongated shapes, we refer to oriented anisotropic spatial filters [16]. Alternatively, spatial Gabor filters [7] can be applied. Combining the spectral filters (4) and the spatial filter (6), we construct the spatio-spectral filter [14] to probe the object's reflectance

$$G_{\lambda^i}^{\sigma_{xy}}(x, y, \lambda) = G_{\lambda^i}(\lambda) * G^{\sigma_{xy}}(x, y). \quad (7)$$

We consider the online measurement of temporal frequency, hence we cannot access information about the "future." Consequently, only a half axis is available: $[-\infty, t_0]$, with t_0 the present moment. Measuring in the domain $t > t_0$ violates causality; a temporal Gaussian filter has infinite extent and, consequently, is only causal over a complete t axis. A reparameterization $s(t)$ of the time axis t is required, such that the filtering of the $s(t)$ domain with a Gaussian is uniform and homogeneous [12]. The requirement of uniform and homogeneous sampling should be independent of the unit of time. Therefore, sampling in the s -reparameterized time axis should be uniform and homogeneous for both clocks t and at , where a is a constant representing a different time scale. Now consider a periodic generator of events, of which the periodicity is estimated in the two time scales $s(t)$ and $s(at)$. Beforehand, no periodicity was more likely than any other. In other words, the probability density function (pdf) f of periodicities as a function of the reparameterized time $s(t)$ in a finite time window is a constant: $f(s(t)) = c$. Further, we require the pdf in the $s(t)$ domain, $f_1(s(t))$, and the pdf in the $s(at)$ domain, $f_2(s(at))$ to be equal: $f_1(\cdot) = f_2(\cdot)$, or, applying the substitution rule when swapping variables: $f_1(s(t)) = (\partial s(at)/\partial s(t)) \cdot f_2(s(at))$. From the latter equation, it follows that the mapping function $s(t)$ must be logarithmic [12]: $f_1(\log(t)) = (\partial \log(at)/\partial \log(t)) \cdot f_2(\log(at)) = (\partial \log(a) + \log(t))/\partial \log(t) \cdot f_2(\log(a) + \log(t)) = f_2(\log(a) + \log(t))$, thus $f_1(\cdot)$ equals $f_2(\cdot)$ except for a shift $\log(a)$. Requiring $f_1(\cdot) \equiv f_2(\cdot)$ implies that both pdfs are




Fig. 1. Sensitivity curves of the spectral probes $G^{\sigma\lambda=55\text{ nm},\lambda_0=540\text{ nm}}(\lambda)$, $G^{\sigma\lambda=55\text{ nm},\lambda_0=540\text{ nm}}(\lambda)$ and $G^{\sigma\lambda=55\text{ nm},\lambda_0=540\text{ nm}}(\lambda)$, approximately colorimetric with the CIE 1964 XYZ sensitivity curves hence with RGB sensitivities [14].

constant. Hence, the reparameterization $s(t) = \log(t)$ satisfies the real-time requirement. Note that we do not single out any position or visible wavelength in the spatio-spectral measurements. In analogy, we do not single out any *range* of time.

For temporal frequency analysis, it is natural to turn to the Fourier domain. With the logarithmic rescaling s of the time dimension t , the Fourier transform of a periodic function in t , $f(t)$ becomes in the s domain: $\int f(t(s)) \cdot (\partial t(s)/\partial s) \cdot e^{-2\pi i t(s)u} ds$, where $t(s)$ is the inverse of $s(t)$, $t(s) = e^s$. Locally weighing the function f with a kernel g , to obtain a joint representation in time and temporal frequency, results in: $\int f(e^s) \cdot g(e^s - e^{s'}) \cdot (\partial e^s/\partial s) \cdot e^{-2\pi i e^s u} ds'$, where g is the logarithmically rescaled Gaussian filter from the t domain. The Fourier transform can thus be rewritten: $\int f(e^s) \cdot G(s - s') \cdot e^s \cdot e^{-2\pi i e^s u} ds'$, with $G(s)$ a Gaussian function. Thus, in the s domain, we get a convolution of the e^s transformed periodic signal f with a kernel $G(s) \cdot e^s \cdot e^{-2\pi i e^s u}$. Translating back to the time domain t , the kernel results in the temporal frequency filter: $G(\log(t)) \cdot t \cdot e^{-2\pi i t u} \cdot (\partial s(t)/\partial t) = G(\log(t)) \cdot e^{-2\pi i t u}$. In full form, we get the temporal frequency filter


$$\tilde{G}^{u_0;\sigma_t,\tau}(t) \equiv \frac{1}{\sqrt{2\pi}\sigma_t} e^{-\frac{\log\left(\frac{t_0-t}{\tau}\right)^2}{2\sigma_t^2}} e^{2\pi i u_0 t} \quad (8)$$

with t_0 the present moment and where τ scales the logarithmic reparameterization hence determines the position of the maximum of the temporal frequency filter. The scaling of the filter determines its extent and is given by σ_t . The shape of the obtained filter resembles auditory temporal frequency filters [10], [17], [18]. Fig. 2 depicts the temporal frequency filter, together with its logarithmically rescaled Gaussian envelope.


The combination of the spatio-spectral filters (7) and the on-line temporal frequency filter (8) yields the online temporal frequency filter for color video

$$\tilde{G}_{\lambda_i}^{\sigma_{xy};u_0;\sigma_t,\tau}(x,y,\lambda,t) \equiv G_{\lambda_i}^{\sigma_{xy}}(x,y,\lambda) * \tilde{G}^{u_0;\sigma_t,\tau}(t). \quad (9)$$

The spatial scale parameter σ_{xy} of the filter determines its spatial extent. Although dependent of the distance between the camera and the object, we will in general consider the object's surface at a coarse scale. The temporal frequency selectivity of the filter depends on the frequency tuning parameter u_0 . We



(a)



(b)

Fig. 2. Reparameterized temporal Gabor component of the temporal frequency filter. Tuning the parameters σ_t and τ determine the shape of the temporal component; we leave the temporal frequency parameter unchanged. Note that (a) has a smaller delay than (b), but a larger temporal extent. The filter parameters are: (a) $t_0 = 0$, $\sigma_t = 5/4$ frames, $\tau = 2$ frames, $u_0 = 1/20$ c/f, and (b) $t_0 = 0$, $\sigma_t = 3/4$ frames, $\tau = 4$ frames, $u_0 = 1/20$ c/f.

do not change the time unit: for the temporal scale parameter σ_t we simply choose $\sigma_t = 1$ frame. As a consequence, the other temporal scale parameter, τ , can be directly related to the tuned temporal frequency. We select the temporal scale a multiplicative of the inverse temporal frequency: $\tau = c(1/u_0)$ frames, with c a constant. As a result, the temporal shape of the filter does not depend on the tuned temporal frequency. Further, the effective temporal extent of the filter directly relates to the temporal frequency, based on the Nyquist theorem that frequency can only be determined if a period of the signal can be resolved.

B. Properties

For color video that is integrally available, the temporal frequency filter is not restricted to a half time axis. As a consequence, the temporal frequency filter does not have to be reparameterized and periodicity can be measured by a temporal Gabor filter [see Fig. 3(b)]. We consider the properties of the offline and online temporal frequency filters, which have different shapes (see Fig. 3; the filters have identical parameters).

The temporal frequency measurement is a local correlation of the periodical color video signal and the temporal Gabor filters. The response of the online filter is asymmetric in time with a fast rise and a slow decay. The online filter hence provides a better fit to an onset of a periodic event in the video data than the offline filter. Fig. 4 demonstrates that the online temporal frequency

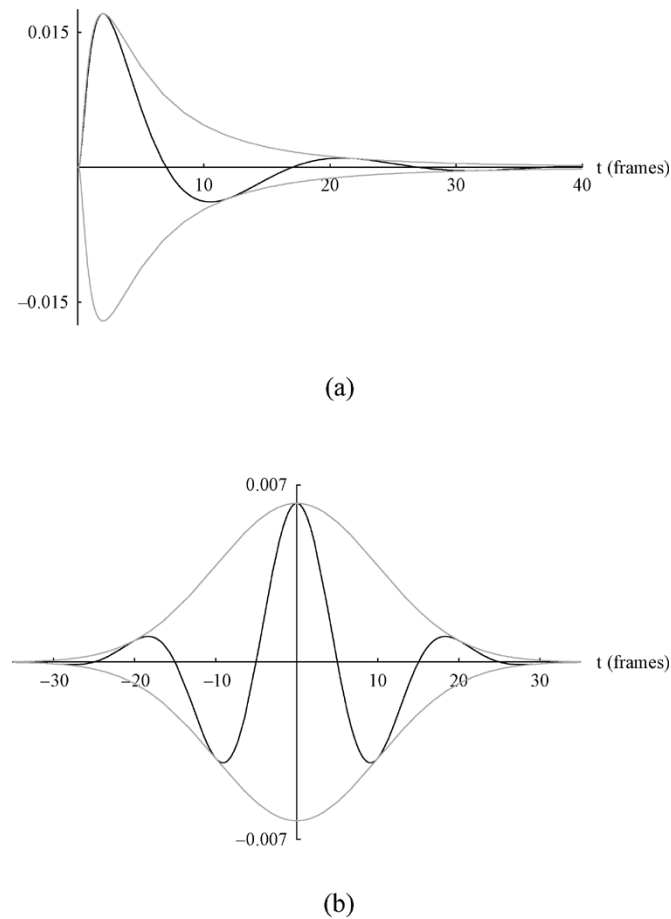


Fig. 3. (a) Online and (b) offline temporal frequency filters for $u_0 = 1/20$ c/f. (a) The time windows of the online and offline filter differ due to the delay $\tau = 2$ frame of the online filter. Note the resemblance in the shapes of the online and offline filter for the past time axis, while the constraint of online filtering is fulfilled. The integral of the filters is normalized to unity, which for the online filter yields a maximum of approximately twice the maximum of the offline filter. Consequently, the online filter will have a faster and higher response than the offline filter.

filters respond faster and decay faster than the offline filters. The online filters respond and decay approximately after one period of the signal plus the delay of the filter τ , see indication at $u_0 = 0.05$ cycles/frame (c/f) $\equiv 20$ frames/period in Fig. 4. In contrast, the offline filters respond and decay approximately after 2.25 periods, being 45 frames at $u_0 = 0.05$ c/f. Hence, the online filter reacts significantly faster than the offline filter.

To determine the temporal frequency selectivity of the filters, we turn to the Fourier domain. See Fig. 5 for the Fourier transforms of the online and offline filter. The online filter is not well localized in the Fourier domain. As a consequence, the online filter yields a low response to higher frequencies than the frequency it is tuned to. However, the Fourier transform of the online filter shows a narrow peak at the tuned frequency. Hence, the online filter is more narrowly tuned to frequencies than the offline filter.

The narrow frequency selectivity of the online filter is demonstrated in Fig. 6. The online filter bank is tuned to dense temporal frequencies. We relate the discrimination quality of the filter bank to the variance of its responses. The variance of the online temporal frequency filter bank is lower than the variance of offline filter responses.

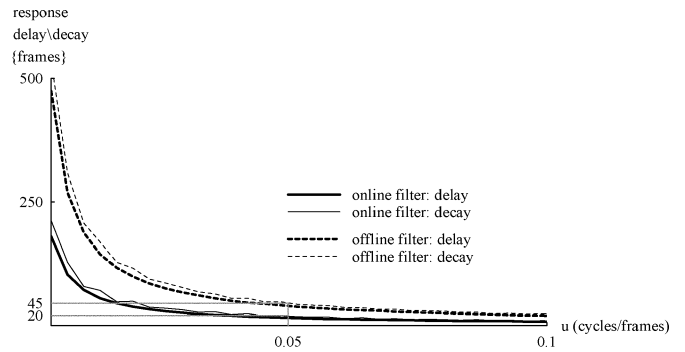


Fig. 4. Response delays (thick lines) and decays (thin lines) of online (solid lines) and offline (dotted lines) temporal frequency filters tuned to signals with various frequencies. The online filters respond and decay approximately after one period of the signal plus the delay of the filter τ , see indication at $u_0 = 0.05$ c/f $\equiv 20$ frames/period. In contrast, the offline filters respond and decay approximately after 2.25 periods, being 45 frames at $u_0 = 0.05$ c/f. Hence, the online filter reacts significantly faster than the offline filter.




Fig. 5. Fourier transforms of the online (a) and offline (b) temporal frequency filters from Fig. 3. Note the narrow peak and the heavy tail of the online filter, compared to the Gaussian shape of the Fourier transformed offline filter.

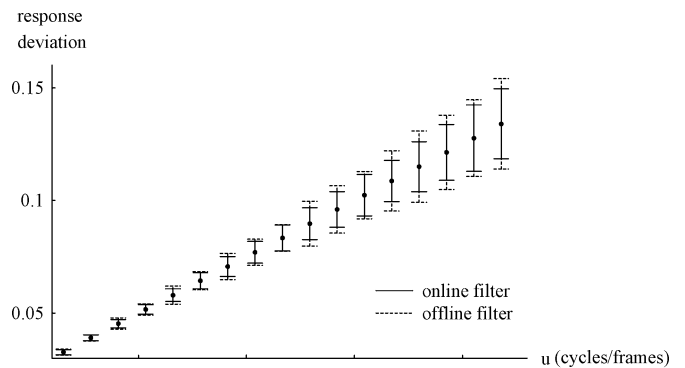


Fig. 6. Frequency selectivity of the online (solid bars) and offline (dotted bars) filter. Frequency selectivity is derived from the variance of the responses of a bank of online and offline filters tuned to dense frequencies (the bars indicate a magnification of 200 times the variance of responses).

We conclude that the online temporal frequency filter achieves higher acuity as it 1) responds and decays faster and 2) can be narrowly tuned to a particular frequency.

C. Algorithm

In the sequel, we define the online temporal frequency measurement for a particular color channel $\hat{E}_{\lambda^i}^{\sigma_{xy}; u_0; \sigma_t, \tau}$ as the magnitude of the complex response of the filter. Filter responses to different color channels are combined by considering their magnitude

$$\hat{E}_{\sigma_{xy}; u_0; \sigma_t, \tau} = \sqrt{\sum_{i=0}^2 \left(\hat{E}_{\lambda^i}^{\sigma_{xy}; u_0; \sigma_t, \tau} \right)^2}. \quad (10)$$

We consider multiple temporal frequency filters, tuned to dense but fixed frequencies, ranging from $1/75$ to $1/7$ cycles per frame. To prune the filter bank, for instance to a range of temporal frequencies that was observed last, a gradient ascent method may be used, taking a filter's response (., its correlation with the signal) as input. Further, the filter is parameterized with a spatial scale. In the experiments, we will preselect a particular spatial scale dependent of the size of the object and the "smoothness" of its motion. Alternatively, the scale of the spatial filter may be derived from scale selection. A common practice is to select the scale according to the maximum of the Laplacian filter [19].

The response of the temporal frequency filter is inherently delayed, depending on the temporal shape of the filter. Responses of filters tuned to lower temporal frequencies are longer delayed. In the experiments, we will both illustrate the delays of different filters and responses where we have aligned filter response delays. The temporal frequency filters primarily respond at half periods of the periodic motion, with alternating magnitudes. We, therefore, integrate the filtering result over a past time window of one period of the filter. Further, we normalize this integration for the size of the time window. We assume the reoccurring surface to have a large spatial extent. Therefore, we spatially pool the responses of the temporal frequency filter. The pooled measurements are thresholded to determine periodicity detection. We identify the frequency as the tuned frequency u_0 of the filter that, after spatially pooling by summation, yields the maximum. As a consequence, we constrain ourselves to the periodic motion of one object. Further, we assume that the maximum spatially pooled response is representative of the periodicity of the object under investigation. Segmenting a frame based on spatially localized responses of temporal frequency filters would overcome the problem of measuring motion periodicity of multiple objects. In the experiments, we will constrain ourselves to demonstrating the robustness of the temporal frequency filter for both stationary and nonstationary single objects moving periodically and quasiperiodically.

III. APPLICATION TO COLOR VIDEO

In this section, we apply a bank of temporal frequency filters to color video of natural scenes. We consider: (a) stable and changing periodic motion of (b) stationary and nonstationary objects with (c) smooth surfaces and regularly textured surfaces (gratings). For all experiments, the color video is recorded by a RGB digital video camera (JVC GR-D72) at 768×576 pixels video frame transfer sampled at 25 frames/s.

A. Periodic Zoological Motion

In this experiment, we *detect* the temporal frequency of the periodic motion of two (stationary) anemones.

Fig. 7 shows a fragment of color video of the periodic motion of a large anemone and a small anemone. The frames are shown in increasing order, from left to right, indicating quarter-periods of the motion of the large anemone. The frames are represented by the three color channels (a) \hat{E} , (b) \hat{E}_λ , and (c) $\hat{E}_{\lambda\lambda}$. The large anemone is located in the center and visible in all color channels (a)–(c). The small anemone is located in the lower left region,




Fig. 7. Color video of the periodic motion of a large anemone and a small anemone. The frames are shown in increasing order, from left to right, indicating quarter-periods of the motion of the large anemone. The frames are represented by the three opponent color channels (a) \hat{E} , (b) \hat{E}_λ , and (c) $\hat{E}_{\lambda\lambda}$. The large anemone is located in the center and visible in all color channels (a)–(c). The small anemone is located in the lower left region, and is only visible in the "green-red" opponent color channel $\hat{E}_{\lambda\lambda}$ [(c), indicated with a circle]. The large anemone moves periodically at $1/18$ cycles per frame, whereas the small anemone moves periodically at $1/8$ cycles per frame. Note that the areas over which the anemones move are marginal, which makes the detection of the anemones' periodic motion nontrivial.

and is only visible in the "green-red" opponent color channel $\hat{E}_{\lambda\lambda}$ [(c), indicated with a circle]. The large anemone moves periodically at $1/18$ cycles per frame, whereas the small anemone moves periodically at $1/8$ cycles per frame. Note that the area over which the anemones move are marginal, which makes the detection of the anemones' periodic motion nontrivial.

We analyzed the frequencies of the anemones' periodic motion at a spatial scale $\sigma_{xy} = 2$ pixels to moderately smooth the signal.

Fig. 8 again depicts a fragment of the color video of the periodic motion of the large anemone and small anemone. The frames cover one period of the motion of the large anemone. The frames are represented by the color channel $\hat{E}_{\lambda\lambda}$ (a), which has most discriminative power. Responses of temporal filters tuned to various frequencies are shown (b)–(g). The filter in (d) is tuned to the frequency of the large anemone. Its response is higher than the responses of filters tuned to temporal frequencies that are slightly lower [(b) and (c)] and higher [(e) and (f)]. We emphasize the inherent delays of the filter: a filter tuned to a lower frequency has a longer delay. The response shown in (d) is higher than the threshold set to determine periodicity. This is




Fig. 8. Color video of the periodic motion of a large anemone and a small anemone. The frames cover one period of the motion of the large anemone. The frames are represented by the color channel $\hat{E}_{\lambda\lambda}$ (a), which has most discriminative power. Responses of temporal filters tuned to various frequencies are shown (b)–(g), where high responses are indicated in white. The filter in (d) is tuned to the frequency of the large anemone. Its response is higher than the responses of filters tuned to temporal frequencies that are slightly lower [(b) and (c)] and higher [(e) and (f)]. The response shown in (d) is higher than the threshold set to determine periodicity. We emphasize the inherent delays of the filter: a filter tuned to a lower frequency has a longer delay. The response of the filter tuned to the periodicity of the small anemone [(g), indicated] is higher than the threshold set to determine periodicity. Despite the isoluminance and weak contrast between the small anemone and its background in the color video fragment (a), the proposed filter strategy was able to detect and identify its periodicity. The frequency parameters are: (b) $u_0 = 1/14$ c/f; (c) $u_0 = 1/16$ c/f; (d) $u_0 = 1/18$ c/f; (e) $u_0 = 1/20$ c/f; (f) $u_0 = 1/22$ c/f; (g) $u_0 = 1/8$ c/f.

also the case for the response of the filter tuned to the periodicity of the small anemone [(g), indicated]. Despite the isoluminance and weak contrast between the small anemone and its background, the proposed filter strategy was able to detect and identify its periodicity. Note that the filter responds to the periodic motion of both the large and small anemone. The ambiguity in the filter's response is caused by the approximate harmonics formed by the frequencies of the motion of the two anemones.

In the description of the algorithm (Section II-C), we mentioned the integration of filter responses over a small time window. As the temporal frequency filter responds maximally at half periods of a periodic event, integration over a half period provides a stable response. In the sequel, we consider integrated responses. For convenience of display, we will align the filter delays with the present moment such that the responses can be compared at single time instances. Fig. 9(a) shows frames

at full periods of the periodic motion of the large anemone. Fig. 9(b) and (c) depict the integrated and aligned responses of the filters tuned to the frequencies of the large and the small anemone, respectively. Integrating the responses of a filter over a half period of the filter provides stability, as demonstrated by the detection of the periodicity of motion in Fig. 9(b) and (c).

The spatial extent over which the two anemones move, pops out from the responses of the filters tuned to their periodic motion. The high responses of the temporal frequency filters evidently reflect the periodicity of the objects under investigation.

B. Periodic Animal Motion

In this experiment, we *identify* the temporal frequency of the periodic motion of a flying bird. The frequency of the bird's wings are a measure of its velocity. Fig. 10 shows a fragment of color video of the periodic motion of a flying bird. The frames




Fig. 9. Color video of the periodic motion of a large anemone and a small anemone. The frames are randomly selected and represented by the color channel (a) $\hat{E}_{\lambda\lambda}$. Two temporal filters [(b) and (c)] are tuned to the respective frequencies of the large and small anemone, respectively, $u_0 = 1/20$ cf and $u_0 = 1/8$ cf . For convenience of display, we have aligned the longer delay (b) and smaller delay (c) of the filter responses with the moment at which the frames were presented. Integrating the responses of a filter over a half period of the filter provides stability over frames within this half period, as demonstrated by the detection of periodicity at randomly selected frames. The spatial extent over which the two anemones move is detected. See “Supplemental Material” for the original color video plus overlaid responses. .

cover one period and are represented by the intensity channel \hat{E} , which contains most discriminative power. Note the variation in the location of the bird.

The bird’s motion inherently causes a translation. To correct for the bird’s translation in subsequent frames, we apply kernel-based tracking with scale adaptivity [20]. We thus exploit a prior model of the bird. For approaches that include automatic motion segmentation we refer to [9], [21], and [22]. The bird’s distance, hence, its perceived size, is normalized by a scaling of the tracked kernel regions, before applying the online temporal frequency filter to the obtained regions. We emphasize that the following experiment’s robustness to, for instance, clutter and occlusion, heavily depends on the tracking of the object. However, tracking objects is not our primary concern here, and, therefore, we will not elaborate on this part of the experiment. Fig. 11 shows the tracking results at half periods of the bird’s motion.

In the frames that differ exactly one period in time (for example, images 1 and 3), the bird has not the same pose. The “misalignment” of the bird’s wings is caused by the low sampling rate compared to the high frequency of its moving wings. Due to the misalignment of the bird’s wings, the problem of identifying the temporal frequency of the bird’s motion is not trivial.



Fig. 10. Color video of the periodic motion of a flying bird. The frames cover one period and are represented by the intensity channels \hat{E} , which contains most discriminative power. Note the variation in the location of the bird.



Fig. 11. Normalized frames as a result of tracking the bird. The frames are taken at half periods of the periodic motion.

In the sequel, we only depict the zeroth order spectral derivative measurement (i.e., the luminance) for display convenience. Further, due to the relatively large spatial extent of the bird, we analyzed the temporal frequency of its surface at a fairly large spatial scale $\sigma_{xy} = 5$ pixels. The advantage of the spatial extent of the filter is that contextual information is incorporated, making the filter robust to the “misaligned” pose of the bird (see “Supplemental Material”).

Fig. 12(a) shows frames at full periods of the flying bird. The temporal frequency of the bird’s periodic motion changes within these samples, i.e., the motion is quasiperiodic. We annotated the temporal frequencies at the full periods. Note that the frequencies differ only one frame per period. At the full periods, we measured and *identified* the temporal frequency of the periodic motion. The responses depicted in Fig. 12(c) and (d) identify alternatively the temporal frequency of the bird’s motion, see the indication. Due to the small differences in the actual frequencies apparent in the bird’s motion, the responses do not differ much. Nonetheless, the identification resembles the annotation. Lower responses of filters tuned to slightly different temporal frequencies are included in Fig. 12(b) and (f).

C. Velocity of Moving Gratings

In this experiment, we measure the *velocity* of a moving grating. A moving grating may be characterized by its orientation, spatial frequency and temporal frequency. The grating velocity is determined by its temporal frequency divided by its spatial frequency [2]. We, therefore, identify both the temporal and spatial frequency.

Analogously to the temporal frequency filter derivation, we analyze spatial frequency in the Fourier domain. When translating back to the spatial domain, we obtain the 2-D spatial Gabor filter [7]

$$\tilde{G}^{\sigma_{xy}, v_0, w_0}(x, y) \equiv G^{\sigma_{xy}}(x, y) e^{2\pi i(v_0 x + w_0 y)}, \quad i^2 = -1 \quad (11)$$

with (v_0, w_0) the frequency in cycles per pixel for two dimensions. The radial center spatial frequency $\sqrt{v_0^2 + w_0^2}$ is given in cycles per pixels and $\tan^{-1}(w_0/v_0)$ represents the orientation of the filter.

We substitute the spatial component of the spatio-spectral filter (7) by the spatial frequency component

$$G_{\lambda^i}^{\sigma_{xy}; v_0, w_0}(x, y, \lambda) = G_{\lambda^i}(\lambda) * \tilde{G}^{\sigma_{xy}, v_0, w_0}(x, y). \quad (12)$$

For a particular color channel, we obtain the spatial frequency measurement $\hat{E}_{\lambda^i}^{\sigma_{xy}; v_0, w_0}$ by considering the magnitude of the




Fig. 12. Color video of the periodic motion of a flying bird. The frames represent full periods of the motion. (a) The frames are represented by the intensity channel \bar{E} . The frequency of the motion changes throughout the represented fragment. We annotated the frequency in a time window around the frames, respectively, $1/7$ cf, $1/8$ cf, $1/7$ cf, $1/8$ cf, and $1/7$ cf. Responses of temporal filters tuned to various frequencies $1/6$ cf, $1/7$ cf, $1/8$ cf, and $1/9$ cf are, respectively, shown in (b)–(e). The filters in (c) and (d) are tuned to two frequencies present in the fragment. At frames 1, 3, and 5, the filter tuned to the annotated frequency of $1/7$ cf responds maximally (c). Its response is slightly higher than the responses of filters tuned to temporal frequencies that are slightly lower (b) and higher [(d) and (e)] (see indications). At frames 2 and 4, the filter tuned to the annotated frequency of $1/8$ cf responds maximally (d). Again, its response is slightly higher than the responses of filters tuned to temporal frequencies that are slightly lower [(b) and (c)] and higher (e) (see indications). For the identification of temporal frequency of the bird’s motion in color video, we refer to “Supplemental Material.”

complex filter response. Note that the spatial frequency measurement does not incorporate time as we consider it at a particular moment.

The color video contains waves propagating through water. Let us define the velocity v of the waves as the ratio of the measured temporal frequency $\hat{E}^{\sigma_{xy};u_0;\sigma_t,\tau}$ and the measured spatial frequency $\hat{E}^{\sigma_{xy};v_0,w_0}(x, y, \lambda)$. Consequently, for a particular location (x_i, y_i) with reflectance λ_i at time t_i , we obtain the velocity

$$v = \frac{\hat{E}^{\sigma_{xy};u_0;\sigma_t,\tau}(x_i, y_i, \lambda_i, t_i)}{\hat{E}^{\sigma_{xy};v_0}(x_i, y_i, \lambda_i)} \quad (13)$$

The propagation of the waves has an orientation of approximately 284° , see the fragment of color video in Fig. 13(a). Therefore, we tuned the spatial frequency filter to an orientation of $\tan^{-1}(w_0/v_0) \equiv 284^\circ$, such that the frequency parameters v_0 and w_0 yield a radial center frequency of $\sqrt{v_0^2 + w_0^2}$ cycles per pixel. Further, we selected a spatial scale of $\sigma_{xy} = 8$ pixels, to cover a sufficient area to robustly measure the occurring frequencies, which are approximately in the range of $\langle 1/14, \dots, 1/8 \rangle$ cycles per pixel. Responses of the oriented spatial frequency filter responses to this grating are shown in Fig. 13(b)–(g). In analogy to temporal frequency identification, the identified frequency corresponds to the frequency of the filter with a maximum spatially pooled response, see indications. For instance, the filter response in Fig. 13(f) to the second frame is higher than filters tuned to slightly different frequencies in Fig. 13(e) and (g). Assigning a maximum response to the first frame is more ambiguous: filter responses in Fig. 13(b)–(d) seem very similar. For the first frame, the algorithm appointed Fig. 13(c) as the maximum response, whereas for the third frame (e) gives the maximum response.

Combining the identified spatial and temporal frequency of these moving gratings, we obtain the velocity of the grating (see Table I for randomly selected frames).

The velocity measurements confirm that the velocity of the grating changes gradually. The spatial frequency of the water, and the temporal frequency of its speed change throughout the video. However, the velocity measurements in random frames reflect that the velocity is more or less stable throughout the video fragment (mean $\mu = 2.0$ and standard deviation $\sigma = 0.8$ pixels per frame, whereas for the whole video fragment $\mu = 1.8$ and $\sigma = 1.0$ pixels per frame).

D. Temporal Frequency as an Attentional Attribute

This final experiment illustrates the periodicity of an object’s motion to be an *attentional attribute*. Motion periodicity, like flicker, is a probable attribute to guide visual attention [4]. Debate exists whether only luminance polarity or both luminance and color polarity draws the attention toward the object [4]. Therefore, we only consider temporal regularity apparent in the luminance channel. Recall that the zeroth order spectral derivative filter measures the luminance.

We analyzed a color video fragment showing both stochastically moving leaves in the wind and one periodically moving leaf. The temporal regularity in the latter leaf guides the attention toward that leaf (see “Supplemental Material”). In Fig. 14, frames $\{0, 5, \dots, 40\}$ depict half periods of the periodically moving leaf. The initial amplitude of the leaf is indicated by dashed vertical lines. Frames 0–40 show 4.5 periods of the moving leaf, thus the leaf moves approximately at a frequency of $1/10$ cycles per frame. We overlaid the response of the temporal frequency filter tuned to $u_0 = 1/10$ c/f. The highlighted regions in Fig. 14 indicate the detection of periodicity. Note that the temporal frequency filter responds well after one period, that is, after frames 0–5 have occurred. Frames 45–55 are included to illustrate the immediate decay in the filter’s response after the leaf starts to move slower and with less amplitude.

TABLE I
VELOCITIES OF THE MOVING GRATINGS

Subsequent frames (random starting point)	150	151	152	153	154	155	156	157	158	159
Spatial frequency (cycles per pixel)	8	6	7	13	10	13	13	10	13	13
Temporal frequency (cycles per frame)	13	6	6	7	7	7	7	7	7	7
Velocity (pixels per frame)	0.6	1.0	1.2	1.9	1.4	1.9	1.9	1.4	1.9	1.9

Random frames	33	36	78	115	132	147	158	162	169	183
Spatial frequency (cycles per pixel)	10	13	6	8	13	8	13	13	10	10
Temporal frequency (cycles per frame)	4	4	5	8	8	3	7	7	3	8
Velocity (pixels per frame)	2.5	3.3	1.2	1.0	1.6	2.6	1.9	1.9	3.3	1.3

Velocity of moving gratings. The moving gratings are taken from color video of waves propagating through water (samples are depicted in Figure 13). Hence, the gratings exhibit a temporal frequency. The ratio of the measured spatial and temporal frequency determines the velocity of the grating [2].




Fig. 13. Color video of waves propagating through water. The frames are randomly selected and represented by the intensity channel \hat{E} (a). The spatial frequency of surface of the water changes throughout the represented fragment. Responses of spatial filters tuned to various frequencies (in cycles per pixel) are shown (b)–(g). In analogy to temporal frequency identification, the identified frequency corresponds to the frequency of the filter with a maximum spatially pooled response, see indications. For instance, the filter response in (f) to the second frame is higher than filters tuned to slightly different frequencies in (e) and (g). Assigning a maximum response to the first frame is more ambiguous: filter responses in (b)–(d) seem very similar. For the first frame, the algorithm appointed (c) as the maximum response, whereas for the third frame (e) gives the maximum response.

Since the leaves themselves only differ in their motion, and not in luminance, color, shape, size, texture, or velocity [4], and the object of attention is not in the center of the color video, we conclude that the attention is only due to the object's temporal regularity. Hence, when temporal regularity is considered in isolation, temporal frequency detection draws the focus of attention.

Real-Time Performance: The filter was applied to color video using a Pentium XEON processor at 2.4 GHz. The computation time of the recursive spatial convolutions described in [23] and [24] is independent of the scale σ_{xy} and relates only to the recursion order of the filter and the dimensions of the video. We set the recursion order to 3. For the European PAL and American NTSC MPEG video standard, the dimensions are: 720×578 pixels at 25 Hz and 720×480 pixels at 30 Hz, respectively. Spatial convolutions for PAL (NTSC) consume 21 (17) ms/frame. Computing three periods spanning 3 s in total, the temporal convolution consumes an additional 18 (15) ms. Total computation thus takes 40 (33) ms, achieving real-time




Fig. 14. Color video of the periodic motion of one leaf, in the midst of stochastically moving leaves. The images represent half periods of the periodic motion at every five frames and are represented by the intensity channel \hat{E} . The subscripts denote the frame offset from the first frame that is displayed. The maximum response among different temporal frequency filters detects the temporal regularity (filter frequency: $u_0 = 1/10$ c/f) and is overlaid as a highlight area. When temporal regularity is considered in isolation, temporal frequency detection draws the focus of attention [4]. Hence, the periodically moving leaf hence guides the attention toward itself. See “Supplemental Material” for the original color video plus overlaid responses.

performance. However, these real-time results were obtained for one filter with a large temporal extent. To meet the real-time requirement for multiple and simultaneous filters as described in Section II-C, parallel computation [25] is required.

IV. CONCLUSIONS AND DISCUSSION

In this paper, we have derived an online, real-time temporal frequency filter. The filter measures in space and wavelength spectrum to estimate the object's surface reflectance. The filter measures temporal frequency to determine the periodicity of the

reoccurrence of the surface. Embedded in the scale-space paradigm, the measurement boils down to a four-dimensional filter, representing a Gaussian filter in the spatio-spectral domain and a Gabor filter in the temporal domain. When measuring online, only the past information can be accessed. We, therefore, have applied a reparameterization of the temporal filter to deal with this constraint. We have introduced color to increase discriminative power to determine the reoccurring surface. Additionally, we introduced spatial extent thereby incorporating local information. We have demonstrated that with moving objects that do not periodically exhibit exactly the same pose, spatial contextual information makes the filter more robust. The constructed online temporal frequency filter measures both frequency identification ('what') and frequency detection ("when").

For simplicity, we have assumed that the spectrum that is reflected from the object does not change under object movement. In general, this assumption does not hold for a moving object. A translation of the object relative to the light source causes primarily shadow and shading deviations. Under the Lambertian reflection model [13], the color video signal $E(x, y, \lambda, t)$ may be decomposed into an intensity component $i(x, y, t)$ and the spectral distribution $e(x, y, \lambda, t)$ representing the color at each location: $E(x, y, \lambda, t) = i(x, y, t)e(x, y, \lambda, t)$. A local normalization of the simultaneous measurement of color and temporal frequency $\hat{E}_{\lambda}^{\sigma_{xy}; u_0; \sigma_t, \tau}$ and $\hat{E}_{\lambda\lambda}^{\sigma_{xy}; u_0; \sigma_t, \tau}$ by the intensity measurement \hat{E} are robust against shadow and shading [14], [26].

In our experiments, we have restricted ourselves to the measurement of temporal frequency of periodic events and the velocity of periodic motion. We demonstrated the general applicability of the proposed filter. Further, we have demonstrated that the online temporal frequency filter is more selective for frequency measurements than the offline filter, as it responds and decays faster.

We have left specialized topics that analyze motion of specific kinds in depth out of consideration. We consequently have not attributed motion-based recognition and gait analysis. The experiments incorporate both the detection and identification of temporal frequency of stationary and nonstationary objects moving periodically and quasiperiodically. In color video, the proposed filter has proven to robustly measure the periodicity of natural motion of objects isoluminant with their background hence only visible in color. The filter has shown to segment the periodically moving object from its background. Although dynamic texture algorithms [6] do not extract explicit frequency information, these algorithms are very efficient in detecting temporal regularity. Hence, dynamic texture segmentation [27] may be useful to determine initially a region of interest to initialize the spatial parameters of the temporal frequency filter. Further, we demonstrated the filter, in combination with a spatial frequency filter, to estimate the velocity of moving gratings well. The estimation of velocity from the periodicity of an object's motion is robust due to its invariance to the object's distance. Although with varying distance the spatial scale of the filter has to be updated by either scale selection or tracking kernel normalization, the frequency of the object does not change. On the contrary, motion estimation from optical flow varies with an object's distance. Further, we illustrated the attentional attribute of periodic motion. Determining the focus of attention is im-

portant as it may detect targets for surveillance video. Finally, we provided examples where periodical events are direct consequences of harmonic mechanical systems in equilibrium. The measurement of an object's periodic motion hence may enable a vision system to estimate parameters of the harmonic mechanism under investigation.

ACKNOWLEDGMENT

The authors would like to thank J. A. Burghouts for collecting the color video used for the experiments (Section III). The authors would also like to thank A. W. M. Smeulders for helpful discussions and comments.

REFERENCES

- [1] F. Cheng, W. J. Christmas, and J. Kittler, "Recognizing human running behavior in sports video sequences," in *Proc. Int. Conf. Pattern Recognition*, vol. 2, 2002, pp. 1017–1020.
- [2] E. P. Simoncelli and D. J. Heeger, "Representing retinal image speed in visual cortex," *Nature*, vol. 4, no. 5, pp. 461–462, 2001.
- [3] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 16, pp. 185–203, 1981.
- [4] J. M. Wolfe and T. S. Horowitz, "What attributes guide the deployment of visual attention and how do they do it?," *Nature Rev. Neurosci.*, vol. 5, pp. 1–7, 2004.
- [5] R. Cutler and L. S. Davis, "Robust real-time periodic motion detection, analysis, and applications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 781–796, Aug. 2000.
- [6] G. Doretto, A. Chiuso, S. Soatto, and Y. N. Wu, "Dynamic textures," *Int. J. Comput. Vis.*, vol. 51, no. 2, pp. 91–109, 2003.
- [7] A. C. Bovik, M. Clark, and W. S. Geisler, "Multichannel texture analysis using localized spatial filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 1, pp. 55–73, Jan. 1990.
- [8] P.-S. Tsai, M. Shah, K. Keiter, and T. Kasparis, "Cyclic motion detection for motion based recognition," *Pattern Recognit.*, vol. 27, no. 12, pp. 1591–1603, 1994.
- [9] R. Polana and R. C. Nelson, "Detection and recognition of periodic, nonrigid motion," *Int. J. Comput. Vis.*, vol. 23, no. 3, pp. 261–282, 1997.
- [10] B. A. Olshausen and K. N. O'Connor, "A new window on sound," *Nature Neurosci.*, vol. 5, no. 4, pp. 292–294, 2002.
- [11] J. J. Koenderink, "The structure of images," *Biol. Cybern.*, vol. 50, pp. 363–370, 1984.
- [12] —, "Scale-time," *Biol. Cybern.*, vol. 58, pp. 159–162, 1988.
- [13] D. Judd and G. Wyszecki, *Color in Business, Science, and Industry*. New York: Wiley, 1975.
- [14] J. M. Geusebroek, R. Boomgaard, A. W. M. Smeulders, and H. Geerts, "Color invariance," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 12, pp. 1338–1350, Dec. 2001.
- [15] E. H. Adelson and J. R. Bergen, "Spatiotemporal energy models for the perception of motion," *J. Opt. Soc. Amer.*, vol. 2, no. 2, pp. 284–299, 1985.
- [16] J. M. Geusebroek, A. W. M. Smeulders, and J. Weijer, "Fast anisotropic Gauss filtering," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 938–943, Aug. 2003.
- [17] E. Boer and H. Jong, "On cochlear encoding: Potentialities and limitations of the reverse-correlation technique," *J. Acoust. Soc. Amer.*, vol. 63, pp. 115–135, 1978.
- [18] T. Irino and R. D. Patterson, "A time-domain, level-dependent auditory filter: The gammachirp," *J. Acoust. Soc. Amer.*, vol. 101, pp. 412–419, 1997.
- [19] T. Lindeberg, "Feature detection with automatic scale selection," *Int. J. Comput. Vis.*, vol. 30, no. 2, pp. 117–154, 1998.
- [20] D. Comaniciu, V. Ramesh, and P. Meer, "Kernel-based object tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 5, pp. 564–575, May 2003.
- [21] Y. Song, L. Goncalves, E. Bernardo, and P. Perona, "Monocular perception of biological motion—detection and labeling," in *Proc. Int. Conf. Computer Vision*, 1999, pp. 805–812.
- [22] H. T. Nguyen and A. W. M. Smeulders, "Fast occluded object tracking by a robust appearance filter," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 8, pp. 1099–1104, Aug. 2004.

- [23] I. T. Young and L. J. van Vliet, "Recursive implementation of the Gaussian filter," *Signal Process.*, vol. 44, no. 2, pp. 139–151, 1995.
- [24] I. T. Young, L. J. van Vliet, and M. van Ginkel, "Recursive gabor filtering," *IEEE Trans. Signal Process.*, vol. 50, no. 11, pp. 2798–2805, Nov. 2002.
- [25] F. J. Seinstra and D. Koelma, "User transparency: A fully sequential programming model for efficient data parallel image processing," *Concurrency Comput.: Pract. Exper.*, vol. 16, pp. 611–644, 2004.
- [26] M. A. Hoang, J. M. Geusebroek, and A. W. M. Smeulders, "Color texture measurement and segmentation," *Signal Process.*, vol. 85, no. 2, pp. 265–275, 2005.
- [27] G. Doretto, D. Cremers, P. Favaro, and S. Soatto, "Dynamic texture segmentation," in *Proc. Int. Conf. Computer Vision*, 2003, pp. 1236–1242.



Gertjan J. Burghouts received the M.Sc. degree in computer science from the University of Twente, Enschede, The Netherlands, in 2002. He is currently pursuing the Ph.D. degree at the Intelligent Systems Lab Amsterdam, University of Amsterdam, Amsterdam, The Netherlands.

His main research interests are in color and texture vision, focusing on the application of invariants for object recognition.



Jan-Mark Geusebroek received the Ph.D. degree in computer sciences from the University of Amsterdam, Amsterdam, The Netherlands, in 2000.

His research interests are in front-end vision, especially color and texture vision. His current research concentrates on computational theories for cognitive vision, based on invariant representations and visual attention. He is Assistant Professor with the Intelligent Systems Lab Amsterdam, University of Amsterdam.