PERGAMON

Engineering Applications of Artificial Intelligence 11 (1998) 257–268

Engineering Applications of
ARTIFICIAL
INTELLIGENCE

Contributed Paper

Robust real-time detection of an underwater pipeline

Primo Zingaretti [a,*], Silvia Maria Zanoli [b]

[a] *Istituto di Informatica, Università di Ancona, via Brecce Bianche, 60131 Ancona, Italy*
[b] *Dipartimento di Elettronica e Automatica, Università di Ancona, via Brecce Bianche, 60131 Ancona, Italy*

Received 1 March 1997; in revised form 1 September 1997

Abstract

Currently, the methods of inspection of underwater structures employ remotely operated vehicles, guided from a support vessel by human operators. The risk of losing concentration calls for the development of an intelligent vision, guidance and control system to support the human activity. The paper presents a robust system for the detection and the real-time tracking of submarine pipelines. An active vision system is proposed to predict changes in the scene, and to direct computational resources to confirm expectations by adapting the processing mode dynamically. The system originates from an image-processing algorithm that was previously developed by the authors to recognise the pipeline in the image plane. The accuracy of this algorithm has been enhanced by exploiting the temporal context in the image sequence. The disturbances on acquired images caused by motion are partially removed by a Kalman filter. The filter proves advantageous in supporting the guidance and control of the ROV, and in making the image-processing module itself more robust. Sequences of underwater images, acquired at a constant sampling frequency from T.V. cameras, are used together with synchronised navigation data to demonstrate the effectiveness of the system.

Keywords: Underwater vision; Remotely operated vehicle (ROV); Image understanding; Object tracking; Vision-based guidance; Active vision; Real-time imaging; Kalman filter

1. Introduction

Oil and gas submarine pipelines need periodic inspections to ascertain their condition, and to prevent damage due to fishing activity, turbulent currents and tidal abrasion. Currently, methods for the inspection of underwater structures are based on remotely operated vehicles (ROVs), guided from a support vessel by human operators. Underwater images, acquired by video cameras mounted on a ROV, are shown on two series of screens, one for the operator who is remotely guiding the ROV, and the second for another operator, who detects anomalies on the pipeline with respect to a standard or a previous situation. Unfortunately, this on-line image analysis constitutes a very demanding task for human operators due to the poor quality of underwater images (e.g. lack of contrast, non-uniform lighting, absorption, scattering and the opacity of the medium), the monotonous character of the scene and the need for constant attention over long periods. The risk of losing concentration may be reduced by lowering the speed of the ROV and allowing more time to analyse a single image. This solution, however, requires an increase in the duration of the mission, with a consequent rise in costs. A more interesting solution consists of supporting the human activity by means of an intelligent vision, guidance and control system. This involves dealing with many problems that still relate to research topics like real-time signal processing, sensor data fusion and world modeling. Moreover, the guidance and control-system responses to scene modifications must be compatible with the speed of the ROV, and synchronisation between the artificial vision system and the control system is necessary. Thus, the proposed intelligent system has to use many artificial intelligence techniques at the same time. In particular, image understanding and object recognition should also include the process of selective acquisition of data in space and time, and an active vision system (Aloimonos et al., 1987; Bajcsy, 1988) should be used to control perception.

* Corresponding author. Fax: +39 71 2204474; e-mail: zinga@inform.unian.it.

0952-1976/98/$19.00
PII: S0952-1976(97)00001-3

Remote inspection based on computer systems, endowed with different degrees of autonomy and artificial intelligence, is widely used in industry, not only to control the quality of production but also to facilitate the maintenance of installations (Byler et al., 1995; Davis, 1990). Visual sensing is reported in the literature in connection with various tasks in submarine robotics. For example, the analysis of image sequences, acquired by a camera installed on a moving underwater vehicle, has been employed for recursive depth estimation (Santos and Sentiero, 1994). Image-recognition techniques have been applied to solve simple navigation and guidance problems, such as the recovery of moored objects (Nguyen *et al.*, 1994). Furthermore, vision-based systems have been proposed to solve the problem of relative motion estimation between an ROV and an offshore structure (Jin et al., 1996; Vaganay and Jouvencel, 1996). The present work proposes a robust real-time system for the detection and tracking of underwater pipelines. The research is framed within a more comprehensive project, carried out by the Italian company Snamprogetti and aimed at enhancing the level of automation in submarine pipeline inspection (Conte et al., 1994). The final goal of the project is the development of an intelligent system for ROV guidance and pipeline inspection. Global architectural aspects of such a system have been discussed by Conte et al. (1994). Position control with visual feedback has been considered in Conte et al. (1996). Processing and understanding of video images during pipeline inspection is a problem analysed by Tascini et al. (1996), while an automatic data-analysis system for reasoning about the pipeline status is taken into account by Conte and Zanoli (1997). The image-processing and detection system in the form presented in Tascini et al. (1996) is well suited to applications such as the one discussed in Conte and Zanoli (1997), but its results are too inaccurate for guidance control. The system described in this paper enhances the one presented in Tascini et al. (1996) by exploiting the temporal context in the image sequence to improve the precision of the computation of the pipeline contours. The disturbances on acquired images caused by motion are also partially removed by cascading a Kalman filter with the image-processing module. The system can predict changes in the scene, and can direct computational resources to confirm expectations, thus making the image-processing module more robust, and making feasible the ROV guidance control.

The paper is organised as follows. In the next section, after a description of some important aspects of image acquisition, the image-processing algorithm for the extraction of pipeline contours from images is briefly reviewed. The outline of the proposed intelligent system, in particular the modeling of the environment to include a cascaded Kalman filter, is described in Section 3, and experimental results and conclusions are presented in Sections 4 and 5, respectively.

2. Pipeline contour detection in the image plane

Periodic submarine pipeline inspections are performed making use of ROVs that carry CCD cameras and other sensors. The success of a survey heavily relies on image understanding. In particular, the detection of the pipe contours represents the first and necessary step in controlling the motion of the ROV along the pipeline. For this purpose attention is focused on top view images (see Fig. 1), which show the pipeline lying on the sea bed and allow the pipe contours to be modeled by two straight lines. Thus,

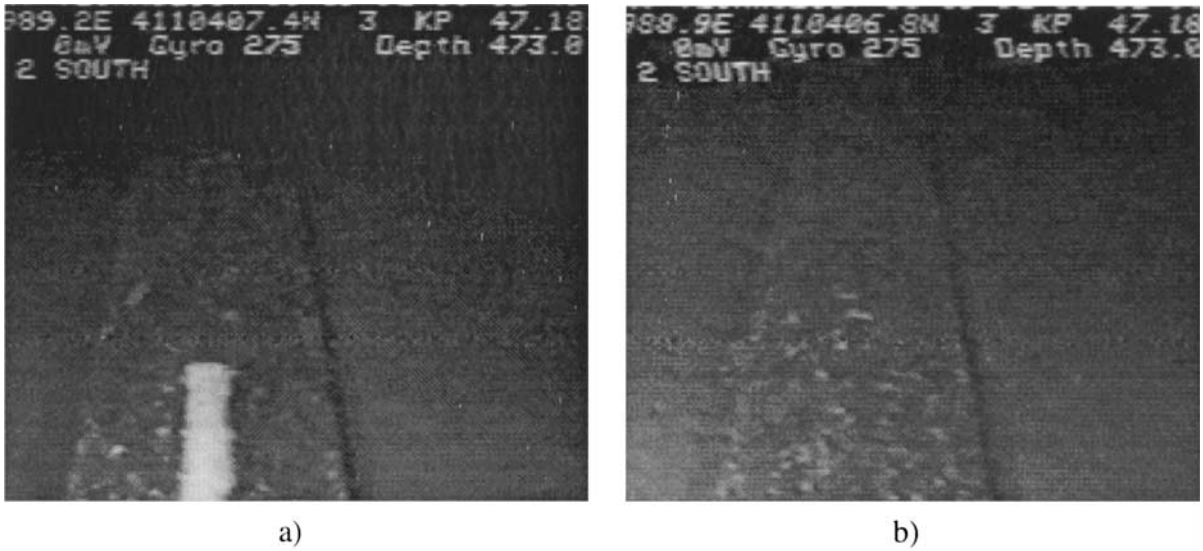

Fig. 1. Underwater images acquired by the top camera mounted on an ROV.

the tracking of the underwater pipeline can be performed directly on the image plane by keeping the pipe in a nearly symmetrical position around image centre. In this Section important aspects concerning image acquisition and the image–processing algorithm for the extraction of pipeline contours are analysed.

2.1. General aspects

In the following paragraphs, the relationships among some important operating parameters (i.e. sampling frequency, ROV speed, image overlap and spatial resolution), the influence of the perspective factor and the effects of motion disturbances in the image plane are described. The acquisition and digitisation processes are performed using a Matrox Image Series processing board (*IM-LC Image Series Manuals,* 1990). The images are digitised with a sampling frequency that guarantees an established percentage of superimposition between two consecutive frames. Image overlap is necessary to increase the robustness of the pipeline contour detection, and to reduce the possibility of failure in the detection of anomalies like cracks in the sheathing, anodes, etc. (Tascini et al., 1996). Once the maximum time required by image-processing operations at each frame is fixed, a trade-off is made between a reduction in the sampling frequency, an increase in the overlap between two consecutive processed frames and an increase in the ROV's speed. Assuming a constant sampling frequency ($1/T$), the speed (V_i) of the ROV may be increased thus reducing the overlap (S_i) between two consecutive sampled frames. This is schematically depicted in Fig. 2, where to the lower speed V_2 corresponds to a greater overlap, S_2. For example, given a vertical image dimension of 450 pixels, a sampling frequency of 1 frame/s and an average spatial resolution of 3 pixel/cm (see below), when the ROV moves at a speed of 1 or 1.5 knots (corresponding to 150 or 225 pixel/s in the image plane) the overlap is, respectively, of two-thirds or one-half of the image.

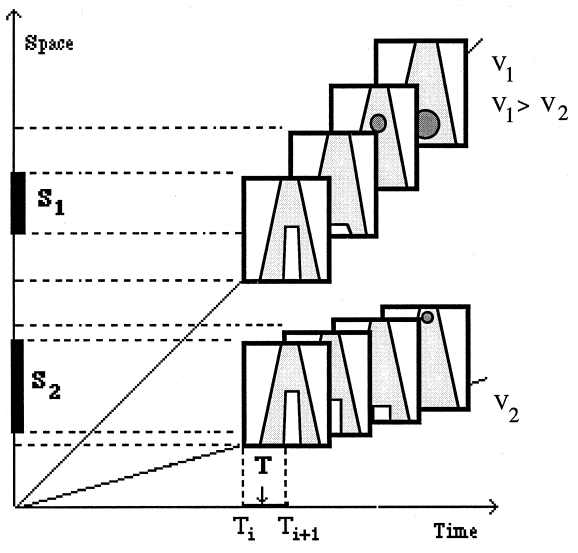

Fig. 2. Relationship between image overlap and sampling frequency.

As mentioned above, the shape of a pipeline can be modeled by a pair of parallel straight lines. However, owing to the influence of perspective, the pipeline contours do not appear parallel in the top-view images. Thus, the model assumed for the pipe consists of two straight lines with small deviations from the upright position, concurrent at the top, in nearly symmetrical positions with respect to the image centre. The known constant value of the pipeline diameter can be used to compute the perspective distortion factor and the spatial resolution at each row in the image. For example, the 60 cm diameter of the pipes we have surveyed here can vary, in the digitised images, from 300 pixels at the bottom to 70 pixels at the top. In particular, an average spatial resolution of 3 pixel/cm may be estimated in the central raster where the pipe diameter is about 180 pixels.

The model adopted for the pipeline enables a simple modeling of the effect of motion disturbances. For inspection purposes, the ROV should track the pipeline at a constant distance and moving at a steady speed. However, actual vehicle motion differs from these constraints causing slight variations in the pipeline's position, orientation and size in the image plane. In particular, a sway linear displacement causes the pipe's contours to sideslip in the image frame [Fig. 3(c)], a heave motion produces a variation in the pipe size in the image [Fig. 3(d)], while pitch and yaw angular displacements affect the slopes of the two straight contour lines [Fig. 3(f) and 3(g)]. Due to the cylindrical shape of the pipe, a limited rolling angle does not produce significant distortions in the image plane [Fig. 3(e)]. Variations in the forward speed (surge motion), on the other hand, do not affect the pipe contour's position in the image plane [Fig. 3(b)]. Given a constant sampling frequency, what is varying in this case is the overlap between two consecutive images (Fig. 2).

2.2. The image-processing algorithm

The algorithm processes video images where the pipe is viewed from the top. The central part of the images, with better lighting, is considered. To overcome the intrinsic non-uniform lighting of underwater images the image is partitioned into horizontal strips (regions), which, being more local, are less affected by this disturbance (mainly present in the vertical direction). A spatial filtering is performed as the first elaboration. Since in the image plane the model of the pipe

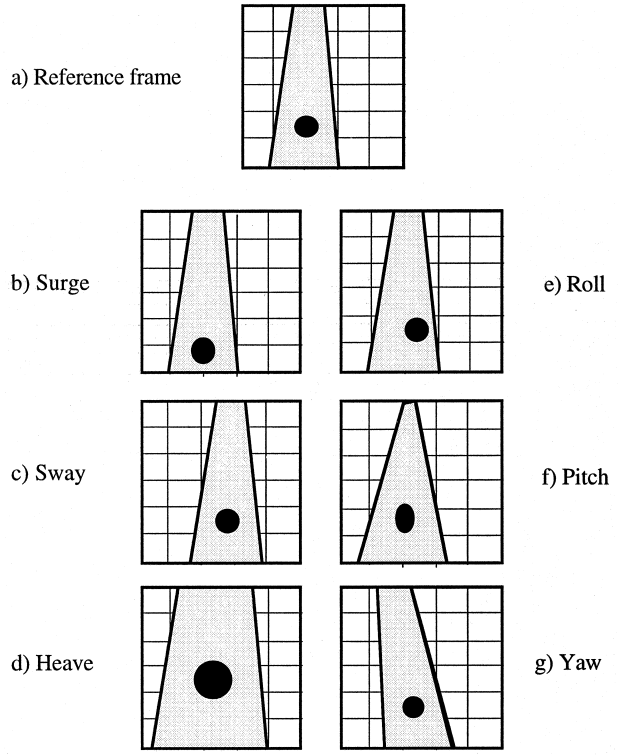

Fig. 3. Variations in position, orientation and size of the pipeline in the image plane. a) Reference image; (b)–(g) effects of motion disturbances on the next digitised frame.

consists of two almost vertical straight lines, a vertical edge detection (Ballard and Brown, 1982) is performed on each acquired image. To overcome the lack of contrast in underwater images, the solution adopted is based on an analysis of the cumulative profiles to select candidate contour points in the edge maps. This profile analysis enables noise filtering to be performed by means of a dense estimation of each edge map. The dense estimation significantly reduces the computational complexity by looking only for clusters of vertical edges, without being misdirected by high-intensity but isolated edge points. The key elements of this density profile analysis can be observed in Fig. 4. Six adjacent horizontal strips with dimensions, in pixels, H and W are selected, starting from the bottom of the edge map. The horizontal profile of each region is computed and stored in a vector of W elements, where each i-element is the sum of the grey levels of the H pixels on the i-column of the region. This procedure, performed by the image-processing board as a primitive function, gathers in the i-element the information contained in H pixels, emphasising the peaks related to the candidate contour points, as well as lowering the casual (not correlated) noisy contributions. Mean value and variance are then computed for each profile. Fig. 4 shows, in the lower part, two of the six profiles of a selected image, i.e. the ones related to the third and the fifth horizontal strips, starting from the top. The number reported on the left of the central line drawn over the profiles denotes the mean value, while the distance of the two extreme lines is related to the variance.

Three tracking procedures, which differ in the way they perform the dense profile analysis, have been developed for the detection of the pipe: *initialisation*, *narrow* and *broad search*. Each procedure has been optimised for a specific assignment. In particular, the initialisation procedure is used in the so-called hooking phase, i.e. in the detection of the pipe when data coming from elaborations on preceding images and relative to the position of the pipe are not available. The narrow-search (NS) procedure is the normal working state. NS is the faster and most robust procedure of the three, as it exploits the information coming from previous image processing. Finally, the broad-search (BS) procedure is used when the NS procedure fails, mainly as a consequence of abrupt horizontal and vertical oscillations during frame grabbing. In addition to the specific performances of each working mode, good throughput in general situations is guaranteed by the

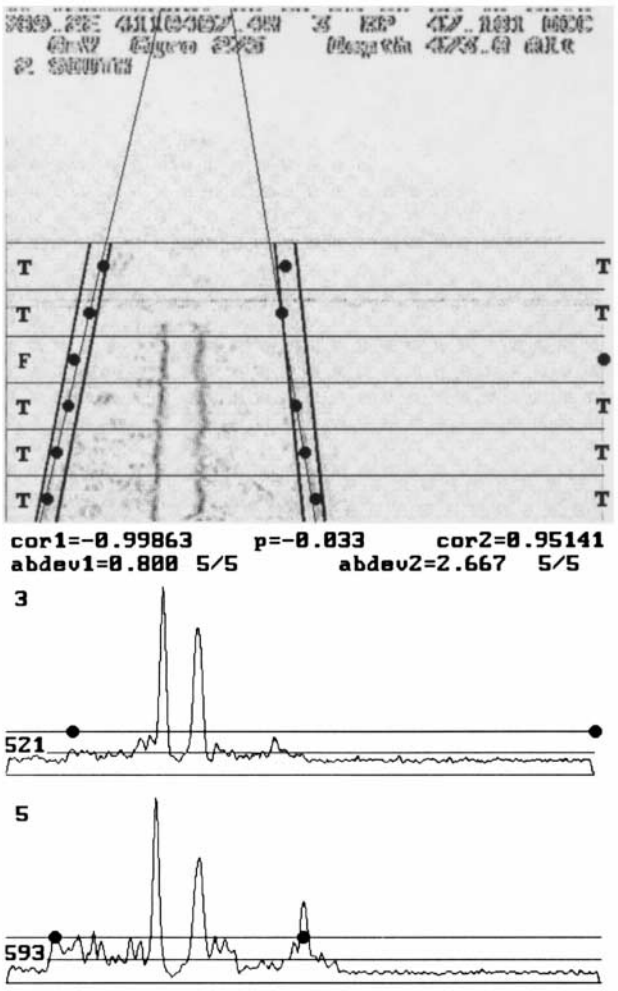

Fig. 4. Profile analysis in narrow-search mode.

architecture of the algorithm, which provides for an effective management of the three procedures.

One of the tracking procedures extracts the two sets of six points that are candidates to represent, respectively, the left and right contour of the pipe. In general, the six points of each group, represented by black circles on the edge map of Fig. 4, are not aligned. For a robust estimation of the corresponding straight lines a robust line-fitting algorithm, implemented by the 'medfit' procedure (Press et al., 1992), is applied. At each iteration step k, the coefficients $a(k)$ and $b(k)$ of a straight line $y = a + bx$ are determined by minimising the squared distances among the six points and the straight line itself. In addition to the coefficients a and b, the average absolute deviation (*abdev*), the number of valid (not discarded) points and the coefficient of linear correlation (*cor*) are computed. In the example reported in Fig. 4, the profile analysis has been performed by the NS procedure. In particular, the right-hand contour has been found with a higher correlation, 0.99 vs 0.95, and a smaller average absolute deviation, 0.80 vs 2.66, than the left-hand one. The letters on either side of each strip denote whether the point selected in the strip has been discarded (F) or not (T) by the 'medfit' procedure.

At start-up, or after an unhooking signal, the initialisation procedure is performed. This procedure selects for each profile the first relative maximum, from both the left- and right-hand sides, exceeding the sum 'mean + variance'. The abscissas of these two selected points are associated to the ordinate middle value of the strip to which the profile refers. The 'medfit' procedure computes the two straight lines that best fit these points. When the resulting straight-line correlation coefficients (*cor1*, *cor2*) are near to 1, and the mean absolute deviations (*abdev*1, *abdev*2) are very low ('good straight lines', GSL), the current frame is considered successfully analysed and the NS procedure is applied to the next frame. Obviously, the presence of other vertical edges to the right or left of the pipe will make the initialisation procedure fail on such frames. The real effect is a delay in the hooking phase, which has to proceed on successive frames. Anyway, the operator can shorten the delay by interactively furnishing two straight lines that approximate the pipeline contours.

In NS mode the algorithm uses two bands of tolerance around the positions of the last GSL, and the analysis of each of the six profiles is performed by determining the relative maximums that lie within these bands. In the example of Fig. 4, the tolerance bands are localised within the two pairs of thick parallel straight lines. In particular, in the profile corresponding to the third horizontal strip, starting from the top, no candidate point has been selected for the right-hand contour. In fact, there is no relative maximum exceeding the profile mean value within the tolerance band. In absence of rapid variations of the pipe position in the image plane, this module succeeds in following the pipeline by updating the GSL whenever the straight lines found fit the candidate points well.

As ROV motion differs from a constant speed and from a straight track, two consecutive frames might not be strictly correlated, and the prediction could fail. If the NS mode fails (i.e. it is not able to determine a sufficient number of valid points within the bands of tolerance, or the straight lines obtained do not satisfy some fixed constraints), the algorithm commutes to BS mode. In BS mode the algorithm considers the peaks of each profile and assigns a weight to each peak on the basis of its width and height. Using these weights, the algorithm looks for pairs of peaks, which will become candidate contour points, using the information on the apparent diameter of the pipe in the frame where the current GSL were found. Again, the straight-line coefficients obtained are analysed and, eventually after the updating of the GSL, the algorithm either returns to NS mode or, if failures in the pipe detection repeat for a fixed number of consecutive frames, signals an unhooking condition. The output of the image-processing algorithm consists of the two pairs of straight-line parameters.

3. The intelligent guidance and inspection system

In the Introduction it was mentioned that this work is part of a research project aimed at the development of an intelligent guidance and inspection system. In such a system, preliminary detection of the pipe is required, both for inspection and for guidance purposes. In the first case, i.e. in checking the status of the pipeline looking for the presence of anomalies, slight errors in the computed contour lines do not influence the inspection process. In contrast, even relatively small errors in the position and orientation of the pipe contours in the image plane and, more seriously, the many variations from one frame to the next, negatively affect the ROV guidance control. In fact, a lack of precision in the position measurement leads to continuous action by the motion controller, and a consequent overload for the thrusters.

The image-processing algorithm described in the previous section may provide inaccurate results, or even fail, as a consequence of two main factors: bad image quality and camera motion. The lack of contrast due to scarce illumination and the presence of suspended particles are typical problems in underwater imaging. The presence of marine growth on the pipe, sea bed settlements, auxiliary structural elements, breaks in the external sheathing of the pipe and alien objects near the pipe are possible causes of badly detectable pipe

contours. In addition, as a result of ROV manoeuvres and of the consequent camera motion, the position and orientation of the pipeline in the image plane may vary greatly with respect to the previous frame. Even if the three search modes for the contour extraction procedure operate correctly, independently of the relative position of the pipe to the image, noisy values of the two parameters $a(k)$ and $b(k)$, characterising each contour line, may result. This occurs mainly as a consequence of misplaced tolerance bands (sway and yaw motion), or a different diameter size (heave motion).

Thus, the insufficient accuracy of the algorithm for pipeline contour detection in the image plane described in the previous section has led to the development of an enhanced vision system, capable of supporting the guidance of an ROV in real time. In particular, an active vision system is proposed, to predict changes in the scene and to direct computational resources to confirm expectations, by exploiting the temporal context in the image sequence and by dynamically adapting the processing mode. The ROV guidance task relies on the reconstruction of the relative position of the ROV (or equivalently of the camera, assumed fixed to it) with respect to the pipeline. The task of the active vision system may be formulated directly in terms of the image plane, as keeping the pipe in a nearly symmetrical position around the image centre.

In the context of active perception (Aloimonos et al., 1987; Bajcsy, 1988) artificial intelligence techniques have been applied to machine-vision problems for the explicit representation of goals and goal-directed perceptual processing. Such vision systems require the modeling of the observer and the world in a synergistic way, and an analysis of the interrelationship between action and perception. In particular, the system described here uses a framework that permits the environment to be modeled in terms of descriptions that are qualitatively different. An elementary model of the world (pipeline) or, better, of its representation in the image plane, is adopted. This enables a simple modeling of the effect of the motion of the observer (camera) on the representation of the world. Finally, world evolution is represented by means of a discrete-time space-state representation which, by including the model of known motion effects, is able to predict future world changes. In the following section, a Kalman filtering technique (Jazwinski, 1970; Ledermann, 1980), developed to reduce the effect of noise on the estimation of these parameters, is described. In addition, the Kalman filter (KF) computes a prediction of the two parameter values in the following sampled $(k + 1)$-th image. The predictive capabilities of the filter speed up the image-processing algorithm and make it more robust.

3.1. Environment modeling

A discrete-time state-space representation of coefficient evolution, frame by frame, is derived in this section. In the discussion that follows, a camera inertial with the ROV is assumed. In addition, the camera and ROV reference systems are made coincident. Thus, each motion deviation perceived by the sensors, and each motion control instruction given by the operator manoeuvring the ROV, refers both to the vehicle and to the camera. The assumptions made, however, do not restrict the validity of the results. A simple reference transformation would apply if the two reference systems did not coincide. Furthermore, when independent camera motions are allowed, simple additional computations are needed to account for the relative motion of the camera and the ROV.

Assuming $\mathrm{X}(k) = [a(k), b(k)]$ as the state vector, the *system equation* that describes the system's dynamic behaviour is as follows:

$$\mathrm{X}(k+1)=\mathbf{F}(k)\mathrm{X}(k)+\mathrm{W}(k)+\mathbf{D}(k)\mathrm{U}(k) \tag{1}$$

where $\boldsymbol{F}$ is the transition matrix, which relates the values of coefficients a and b at time k to time $k + 1$. For inspection purposes, the ROV has to follow a straight path, constantly tracking the pipeline. Thus, in normal conditions, the position of pipeline contours in the image plane is not expected to vary from one frame to the next. The transition matrix $\mathbf{F}$ thereby takes on a particularly simple form, i.e. $\mathbf{F}$ is equal to the identity matrix. The vector $\mathrm{U}(k)$ represents an exogenous input, which accounts for deviation of the vehicle track from a straight path. More precisely, known perturbations of the lateral position (Δy) of the vehicle, that is, of the camera, with respect to the centre line of the pipe, and/or perturbations of the orientation ($\Delta\varphi$) in the horizontal plane, are considered. Roll and pitch angular perturbations are in general limited, and their effects have been modeled by the system noise vector $\mathrm{W}(k)$. Heave motion is assumed to be limited as well. As a matter of fact, this assumption is not satisfied during the entire survey process, but since variations in the vehicle depth are, in general, single events, it can be piecewise verified. A future enhancement of the system will include these motions, taken from readings of the altimeter sensor, in the U vector of known system perturbations.

Under the hypotheses that the ROV motion occurs on planes parallel to the plane in which the pipe lies, and that the depth Z_t of the vehicle, i.e. of the camera, is known and constant (angles of pitch and roll negligible), the equations that correlate different camera positions with straight-line parameters a and b in the image plane are (Conte et al., 1996):

$$a = \frac{-\lambda(B\sin\alpha + A\sin\varphi\cos\alpha)}{(B\cos\alpha + A\sin\varphi\sin\alpha)} \tag{2}$$

$$b = \frac{A\cos\varphi}{(B\cos\alpha + A\sin\varphi\sin\alpha)} \tag{3}$$

where the angle α is the inclination of the camera with respect to the normal axis of the pipe (supposed known and constant), λ represents the camera's focal length, and the φ angle is the orientation of a reference frame centred on the camera with respect to that centred on the pipe. Parameters A and B in Eqs. (2) and (3) depend on the pipe radius R, on the lateral displacement Y_t and on the depth Z_t. Their expressions are:

$$A = Z_t - K_j\frac{K_jZ_t + Y_t}{K_j^2 + 1} \tag{4}$$

$$B = Y_t - \frac{K_jZ_t + Y_t}{K_j^2 + 1} \tag{5}$$

where K_j is equal to:

$$K_j = \frac{-Y_tZ_t \pm R\sqrt{Y_t^2 + (Z_t^2 - R^2)}}{Z_t^2 - R^2}. \tag{6}$$

Hence, the matrix **D**, which characterises the parameter variation with respect to known (from navigation sensors) perturbations Δy and $\Delta\varphi$, has the form:

$$\mathbf{D}(k) = \begin{bmatrix} \dfrac{\partial a}{\partial\varphi} & \dfrac{\partial a}{\partial y} \\ \dfrac{\partial b}{\partial\varphi} & \dfrac{\partial b}{\partial y} \end{bmatrix}_{\hat{a}(k),\hat{b}(k)}. \tag{7}$$

To complete the state-space representation, the output of the image-processing module (IP module in Fig. 5) needs to be related to the state variable. This is done by the following *observation equation*:

$$\mathrm{Z}(k) = \mathbf{H}(k)\mathrm{X}(k) + \mathrm{V}(k) \tag{8}$$

where the observation matrix **H** is simply constituted of unit elements, and V(k) represents the measurement noise.

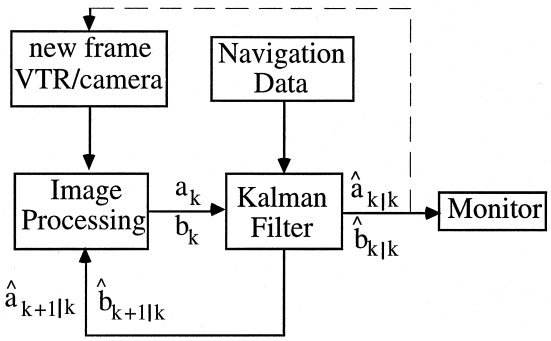

Fig. 5. Architecture of the pipeline-detection module.

Given the state-space representation of Eqs. (1) and (8), a KF that recursively estimates the state of the system is derived. System noise W(k), due to unknown camera motion, and observation noise V(k), resulting from the contour-extraction procedure, are modeled as Gaussian and white, mutually uncorrelated and independent from time to time. Under these hypotheses, the KF is optimal in the sense of the minimisation of the mean square error. For the state estimate at time k, denoted as $\hat{\mathrm{X}}(k|k)$, the following expression holds:

$$\hat{\mathrm{X}}(k|k) = \hat{\mathrm{X}}(k|k-1) + \mathbf{K}(k)[\mathrm{Z}(k) - \mathbf{H}(k)\hat{\mathrm{X}}(k|k-1)] \tag{9}$$

where $\hat{\mathrm{X}}(k|k-1)$ denotes the prediction of the state value at time k, based on the information gathered up to time $k-1$, and K(k) is the Kalman Gain matrix (Ledermann, 1980).

Predictions are given by:

$$\hat{\mathrm{X}}(k+1|k) = \mathbf{F}(k)\hat{\mathrm{X}}(k|k) + \mathbf{D}(k)\mathrm{U}(k). \tag{10}$$

3.2. The system architecture

The filtering process described in the previous section is separately applied to each contour line. At each step k, Eqs. (9) and (10) return both an estimation and a prediction of the two parameters a and b characterising each contour line. The overall system architecture is shown in Fig. 5. The KF module receives as input the pairs (a, b) of contour-line parameters resulting from the image-processing module (see Section 2.2). The navigation data module accounts for known perturbations included in the system by the U(k) vector [Eq. (1)].

Predicted values $\hat{a}_{k+1|k}$,$\hat{b}_{k+1|k}$ are fed back to the image-processing module, while filtered data $\hat{a}_{k|k}$,$\hat{b}_{k|k}$ constitute the final output of the system. These data are displayed on the screen monitor in the form of two straight lines, superimposed on the currently analysed image [Fig. 6(c), 6(f)].

The feedback action performed by the KF enables a reduction and/or a better localisation of the tolerance bands, thus improving the performance in the search for candidate contour points. In fact, the better localisation of the tolerance bands is a direct consequence of the availability of noise-free data which, furthermore, combine the effects of known disturbances. Indirectly, this implies a longer retention of the system in the NS mode, even with a reduced tolerance band. In addition, the contour–detection procedure takes advantage of the estimated parameter values by reducing the number of failures, and the consequent calls to the three procedures.

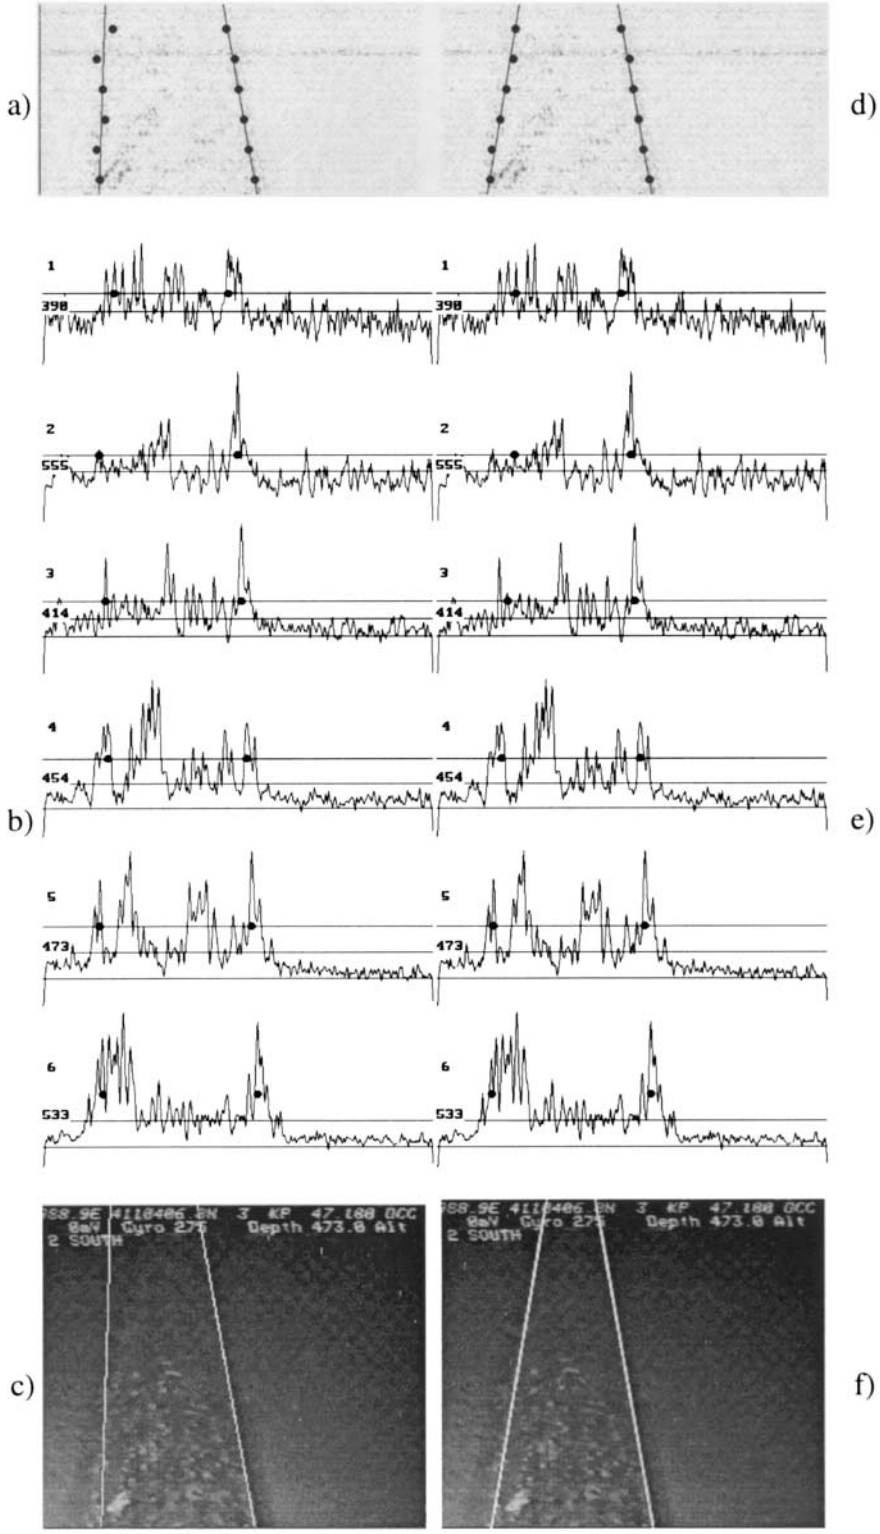

Fig. 6. Results of dense profile analysis and contour straight line estimation: (a)–(c) without Kalman filter; (d)–(f) with Kalman filter.

The whole system, consisting of the KF cascaded with the image processing module, turns out to be computationally quite convenient since the number of implemented filters is limited. More effective noise reduction could in principle be obtained by filtering the information at an earlier stage, namely when points belonging to the pipe contour are selected. This, however, would imply the use of a greater number of filters. Another advantage of this structure is that it gains in flexibility by separating the action of the filter

from the contour-extraction process. The contour-extraction procedure could be modified in accordance with the available hardware and the characteristics of the acquisition process without influencing, in this solution, the filtering stage.

The output of the system consists of the final restitution of the two-straight line equations, and their visualisation over the analysed image. Moreover, the active vision system has been developed with the respect of all the constraints imposed by the guidance control module. Thus if the feedback action represented by a dashed line in Fig. 5 were made operative, visual guidance of the ROV could be realised.

4. Results

The effects of the introduction of the KF module in the system are demonstrated in this section. In particular, the candidate contour points selected by the profile-analysis procedure, and the two resulting pairs of contour straight lines, with or without the KF, are compared. The effectiveness of the cascaded filter on sequences of 1000 frames is also discussed.

The basic processing of the contour-detection procedure, i.e. dense profile analysis and contour straight-line estimation, can be traced in Fig. 6. The edge map resulting from the spatial filtering of the image in Fig. 1(b) is shown in the background of Fig. 6(a) and 6(d). Even if the same edge map is processed by the contour-detection procedure, the resulting black points and straight lines are different in the cases where Kalman filtering is applied [Fig. 6(d)], or not [Fig. 6(a)]. In fact, in the two cases the pairs of points selected for each profile may even be different as a consequence of the bands of tolerance being centred on different positions, that is, based either on the most recently found GSL, or on the straight lines predicted by the KF. Moreover, different tracking procedures and, consequently, bands of tolerance of different width, may be activated, depending on whether or not the filtering process is applied. The horizontal dense profiles relative to each strip are depicted in Fig. 6(b) and 6(e), where the two black circles in each profile correspond to the pair of selected points. In particular, the profiles of Fig. 6(b) were analysed in BS mode, while those of Fig. 6(e) were done in NS mode. In the profiles relative to strips 1, 2, 3 and 6 the different selections of the left point should be pointed out. Consequently, the application of the 'medfit' algorithm gives different contour lines, with different linear correlation coefficients. As can also be visually verified by the better line fitting in Fig. 6(d), the correlation coefficient is higher when the filter is active. Finally, in Fig. 6(c) and 6(f) the resulting straight lines are superimposed on the original grey-level image. It should be pointed out that even if the result in Fig. 6(c) is clearly wrong, the previous image-processing algorithm could nevertheless not discard it, being physically possible, for example as a consequence of a combined yaw and heave motion. In contrast, the pipeline detection system including the KF is more reliable because it does not require additional verification, and can justify abrupt parameter variations only if they have been predicted. In the case of Fig. 6 there were no yaw or heave motions, and the better performance of the system with the embedded KF is due to the different tracking procedure used to analyse the frame. In fact, when contours are not well defined, like the left-hand contour in Fig. 6, many peaks of similar weight may result within the tolerance bands, particularly in BS mode where wider bands are taken. In contrast, the well-defined right contour in Fig. 6 has no competitors, so even a wider band does not influence the candidate-point selection. As the application of the KF allows the frame under examination to be processed in NS mode, the selection of left-hand points may be performed by using a more tightly focused search, by which noisy peaks are excluded even if they have a higher weight.

The effect of a sway motion between successive frames can be traced in Fig. 7, where a sequence of four frames and the relative processing results are shown. In consequence of the previous elaborations, frame (a) is processed in BS mode when the KF is not applied, as can be visually noticed by the absence of the narrow tolerance bands (central column in Fig. 7).

In any case, the contour-detection procedure gives the same results as when using the KF (right-hand column in Fig. 7). In addition, GSL are found so that frame (b) can be processed in NS mode.

The presence of noisy peaks within the tolerance bands leads to the selection of candidate points that result in acceptable straight lines ($cor1 = 0.95$, $cor2 = 0.97$), but do not account for the left-shift of the pipeline in the image plane as consequence of the sway motion. This does not occur when using the KF, where GSL are obtained ($cor1 = 0.99$, $cor2 = 0.98$). The further left-sway motion in frame (c) does not interrupt the NS mode when the KF is active, but it causes a failure in the other case, thus requiring a call to the BS procedure. Again, owing to the availability of good information about the pipe diameter, the BS procedure gives the same results as when using the KF. In addition to the double processing on the same frame, the switch to BS mode produces a further inaccuracy in the extraction of the left-hand straight line of the succeeding frame (d) ($cor1 = 0.98$), with respect to the corresponding one, obtained using the KF ($cor1 = 0.99$). Summing up, the four-frame sequence required only 4 NS procedure calls when using the KF vs 7 (=4 NS + 3 BS) tracking procedure calls in the

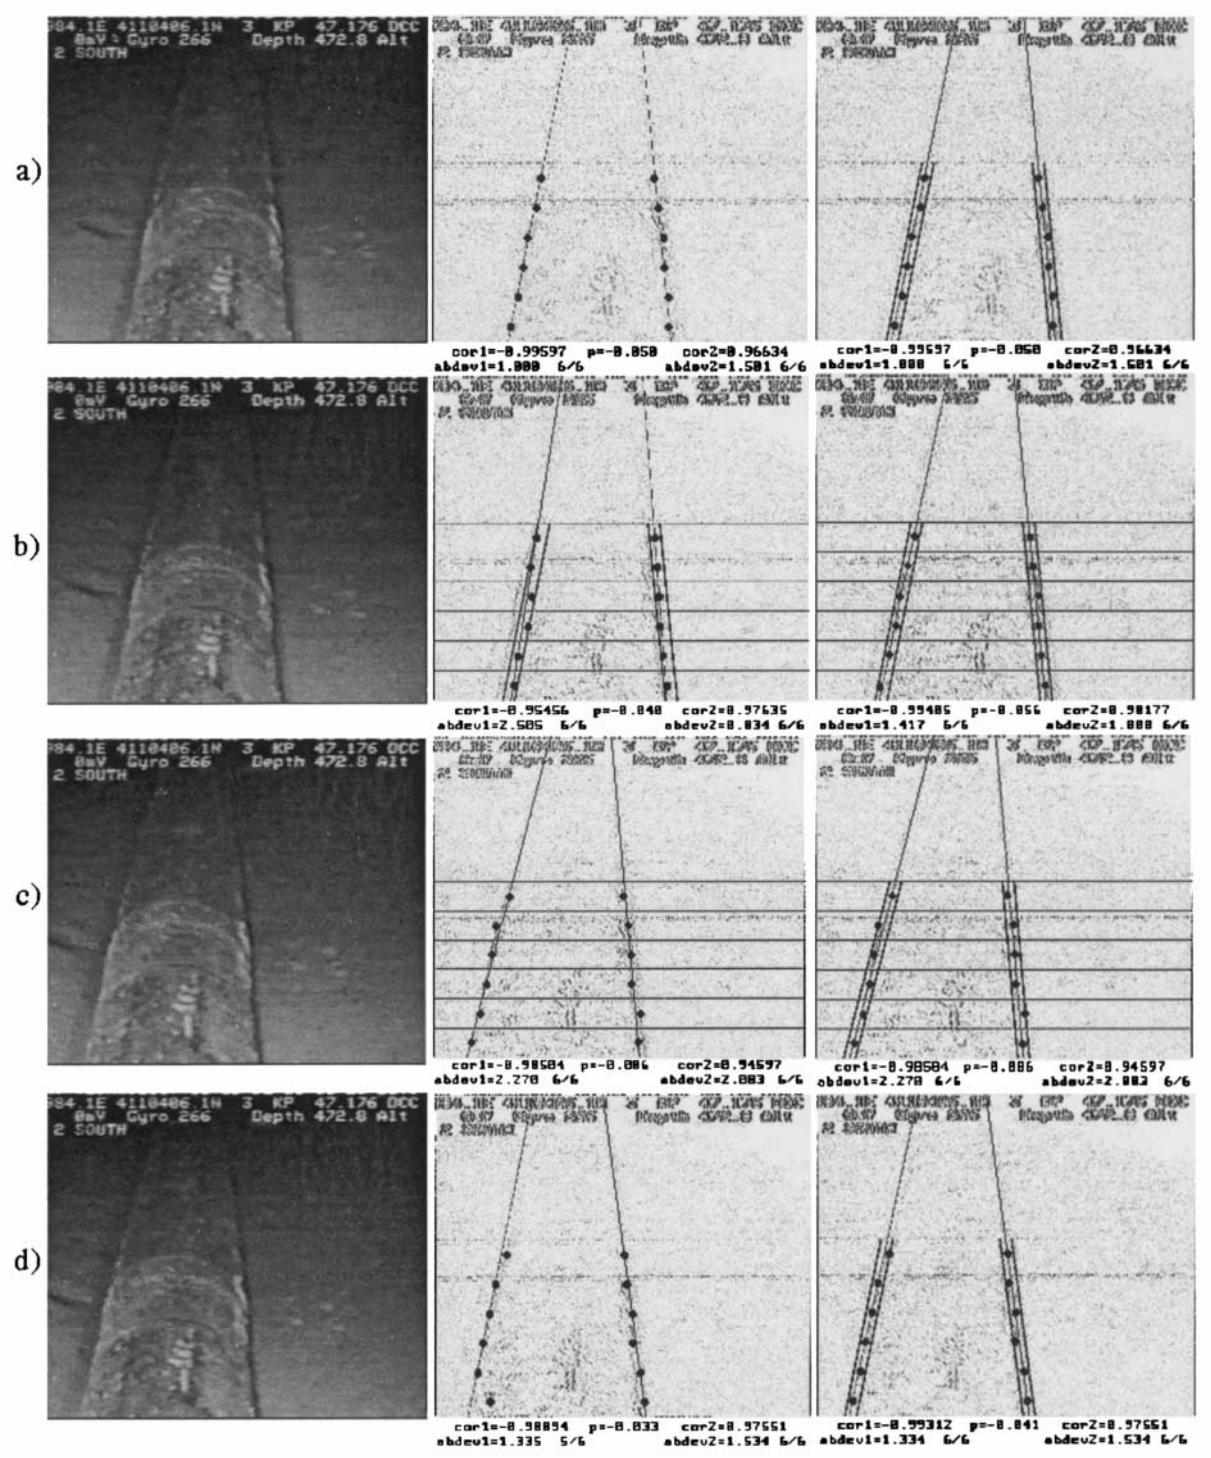

Fig. 7. Processing of a frame sequence. Results of pipeline contour extraction for the image on the left with (right column) or without (central column) the Kalman filter.

other case. In addition, less-accurate straight lines were found without the KF.

To test the performance of this active vision system using the only available pre-recorded video image sequences, it was assumed that the operator controlling the camera behaves as an active vision system, at least in keeping the pipe within each frame. Fortunately, this is the case in all the pre-recorded image sequences. The effectiveness of the cascaded filter was tested in the laboratory on video-recorded sequences. 1000 frames for each sequence were digitised at a constant sampling frequency and stored, in order to compare the performance of the system when the filter is applied and when it is not. Synchronised motion data

Table 1
Comparison of procedure calls during the processing of a 1000-frame sequence, with and without the application of the Kalman filter

Procedure	No filter			Kalman filter		
Re-initialisation	133	18	115	139	9	130
Narrow-search	897	495	402	971	884	87
Broad-search	520	370	150	142	45	97
Total	1550	883	667	1252	938	314
	Total	OK	BAD	Total	OK	BAD

were extracted from available navigation data. A first advantage of the filter application is the reduction of the overall processing time, as a consequence of both the lower number of procedure calls (e.g. 1252 vs 1550 in the sequence analysed in Table 1) and the greater percentage of calls to the NS procedure, which is the fastest mode. In fact, these two factors greatly compensate for the small overload of the KF. A second important advantage is the increased number of successfully processed frames due to the prediction of the filter, which allows for a better focused search. In the example examined in Table 1 the successfully processed frames were 93.8% vs 88.3% in the absence of the filter. An average increment of about 5% was also obtained for the other sequences.

The system, implemented on a PC equipped with an Intel-Pentium CPU and a Matrox Image-LC image processor, is able to detect and track the pipeline, working in real time at a maximum frame rate of about 7.4 frame/s, using images with a vertical dimension of 450 pixels. Thus, a simulated ROV speed of 5.5 kn can be reached, with an image overlap of almost 3/4, or an ROV speed of 7.4 kn if 2/3 image overlap is acceptable. This is a speed that is much higher than that normally reached in inspection missions; consequently, by adopting a lower speed, the extra processing time may be used for solving other tasks involved with pipeline inspection, such as the detection of cracks in the sheathing of the pipe.

5. Conclusions

The aim of this work was to contribute to showing the feasibility of an automatic real-time vision system for supporting human operators in pipeline inspection and ROV guidance. The proposed intelligent system uses very simple geometric models (straight lines) for reducing the computational cost, and for obtaining fast and simple computational procedures. An image-processing algorithm, exploiting the information coming from the processing of the previous frames to extract the parameters characterising the pipeline, has already provided good performance (Tascini et al., 1996). The overall robustness of this algorithm has been increased in the integrated vision system presented in this paper. The effects of motion on the image-acquisition process, that is, the consequent displacements of the straight lines in the image plane, have been included in the equations of a KF cascaded with the image-processing module. The feedback action of the cascaded KF improves the reliability of predictions, and provides a more accurate localisation of the ROV relative to the underwater pipeline. Thus, the active vision system can better support the guidance of the ROV, so as to keep the pipe in a nearly symmetrical position around the image centre. Performance evaluation, accomplished by laboratory simulations using navigation data and images acquired during underwater surveys, has proved the improvements obtained over a previous system.

References

Aloimonos, J.Y., Weiss, I., Bandopadhay, A., 1987. Active vision. Internation. Journal on Computer Vision , 333–356.

Bajcsy, R., 1988. Active perception. IEEE Proceedings 76, 996–1006.

Ballard, D., Brown, C., 1982. Computer Vision. Prentice–Hall, Englewood Cliffs' N.J.

Byler, E., Chun, W., Hoff, W., Layne, D., 1995. Autonomous hazardous waste drum inspection vehicle. IEEE Robotics and Automation Magazine 2.

Conte, G., Zanoli, S.M., 1997. A system theoretic approach to automatic data analysis for inspection of underwater structures. International Journal of Systems Science 28 (8), 737–748.

Conte, G., Zanoli, S.M., Perdon, A., Radicioni, A., 1994. A system for the automatic survey of underwater structures. Proceedings of Ocean 94—OSATES 94 II, 92–95.

Conte, G., Corradini, M.L., Orlando, G., Zingaretti, P., 1996. Discrete-time VSC with visual feedback applied to the position control of an underwater ROV: an emulation study. In: Proceedings of 6th IARP workshop on Underwater Robotics. Touloun, France.

Davis, R.S., 1990. Remote visual inspection in the nuclear, pipeline and underwater industries. Materials Evaluation 48, 797–803.

IM-LC Image Series Manuals, 1990. Matrox Electronic Systems. Dorval, Quebec, Canada.

Jazwinski, A.H., 1970. Stochastic Process and Filtering Theory. Academic Press, New York.

Jin, L., Xu, X., Negahdaripour, S., Tsukamoto, C., Yuh, J., 1996. A real-time vision-based stationkeeping system for underwater robotics applications. IEEE Proceedings of Oceans 96.

Ledermann, W., 1980. Kalman Filtering. Handbook of Applicable Mathematics, Statistics, Vol. IV. Lloyd, New York.

Nguyen, H.G., Kaomea, P.K., Heckman, P.J. Jr, 1988. Machine visual guidance for an autonomous undersea submersible. Underwater Imaging, 980. SPIE, pp. 82–89.

Press, W., Teukolsky, S., Vetterling, W., Flannery, B., 1992. Numerical Recipes in C. 2nd. Cambridge University Press, New York.

Santos, Victor J., Sentiero, J., 1994. The role of vision for underwater vehicles. In: Proceedings of 1994 Symposium on Autonomous Underwater Vehicle Technology. Cambridge, MA, pp. 28–35.

Tascini, G., Zingaretti, P., Conte, G., 1996. Real-time inspection by submarine images. Journal of Electronic Imaging, 5, 4. SPIE-IS&T, pp. 432–442.

Vaganay, J., Jouvencel, B., 1996. Motion estimation for ROV stabilization with a light-stripe sensor. IEEE Proceedings of Oceans 96.

Primo Zingaretti received a Doctoral degree in electronic engineering from the University of Ancona, Italy, in 1984. From 1987 to 1990 he has obtained annual grants at INRCA of Ancona, Italy, for developing medical image-processing systems and artificial intelligence software applications. In 1990 he joined the faculty at the University of Ancona, Italy where he currently holds the position of Research and Development Engineer at the Computer Science Institute. His current research interests concern image processing, image understanding, perception and learning. He is a member of IEEE, IAPR and AI*IA (Italian Association for Artificial Intelligence).

Silvia Maria Zanoli received an M.S. degree in electronic engineering with the maximum score in 1991, and a Ph.D. degree in 1995 from the University of Ancona. Since 1993 she has been working as tutor in the System Theory course at the graduate school in Fermo (Ascoli Piceno). She was a visiting researcher at the MSEL of the Northeastern University, Boston, in 1994/95. At present she is a research fellow at the Department of Electronics and Automation of the University of Ancona. Her research interests concern mobile and underwater robotics, autonomous systems, intelligent control, and automatic systems for data analysis.