JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 20, No. 6, November–December 1997

Using Probability Estimates to Identify Environmental Features for a Nonholonomic Control System

John-David Yoder,* Steven B. Skaar,† and Humberto Arriola‡
University of Notre Dame, Notre Dame, Indiana 46556

A third-order set of nonlinear, ordinary differential equations models the relationship between internally measurable wheel rotations and the position and orientation of an automatically guided vehicle, but these relationships are imprecise, growing increasingly inadequate as their integrals, and the vehicle, proceed from point of departure. An extended Kalman filter (EKF) is used to combine video observations of features on that portion of the environment that does not move, together with the sensed wheel rotations, to produce the ongoing estimates needed for navigation. The experimental usefulness is examined of a byproduct of the filter, the estimate error covariance matrix, to an integrally related process: the process of identifying video observations with features of known location within the environment; these identities are required for application of new vision observations to the state estimates. The goodness of the EKF's probability density functions is experimentally examined by comparing them against actual, accumulated data; experimental results are presented from the use of an extensive theoretical development that assesses, based on relative probabilities inferred from these distributions, the identities of densely occurring, nondistinct cues.

Introduction

NAVIGATION of a mobile, nonholonomic robot through a workspace, such as a home or factory floor, can entail issues interestingly similar in certain respects and dissimilar in others, compared to the autonomous missile- or aircraft-navigation problem. Of particular interest is the dual use made of the estimation error covariance matrix of the extended Kalman filter (EKF)[1] as it is applied to vehicle position/orientation estimates using onboard image-sensor observations of cues with known locations in a constructed indoor setting. The error covariance matrix and state estimates are used here in the usual capacity, in this case in the production of ongoing vehicle-position and vehicle-orientation estimates. This paper discusses the way in which the error covariance matrix $[P]$ can also be used to provide useful probabilities regarding the identities of visual cues, which are used in subsequent EKF-based estimate updates.

Certain kinds of features, such as vertical edges, are common within an artificial environment and relatively easily detected in the image plane; they may also, however, be (practically) indistinguishable from one another using real-time imaging, dense, and frequently not part of the a priori map of the environment. The usefulness of the aforementioned probabilities in distinguishing among many similar visual features is illustrated and reported using special-purpose elliptically shaped wall marks.

The usefulness of cue-identity probability figures determined using output from the EKF depends on the reliability of the estimate-covariance-based joint probability density functions themselves. Because the EKF computations utilize algorithms that are based on linearization, as well as assumptions of uncorrelated (over time) white noise corrupting both the dynamic model as well as the individual measurements, it is prudent, in a given application, to have some empirical validation of the probability distributions predicted by the EKF; it is also prudent to err on the conservative side in utilizing these figures to make cue-identity inferences by using incoming visual data in a way that makes possible very high thresholds of probability for acceptance of a particular cue identity. Also discussed is how both of these provisions are accommodated with a particular mobile-robot system, an automatically guided wheelchair using onboard video observations of wall-mounted features, as well as ways in which the probabilities assigned to cue identities can result in increasingly high levels of certainty by utilizing 1) multiple properties of a single detected feature and 2) multiple features, or cues, appearing in a single observation.

Received Nov. 21, 1996; revision received July 11, 1997; accepted for publication July 21, 1997.

*Graduate Research Assistant, Department of Aerospace and Mechanical Engineering; currently President, Yoder Software, Inc., 3100 Benham Avenue, Elkhart, IN 46517.

†Professor, Department of Aerospace and Mechanical Engineering. Associate Fellow AIAA.

‡Graduate Research Assistant, Department of Aerospace and Mechanical Engineering.

Kinematic Differential Equations of Motion

The particular nonholonomic kinematics of the wheelchair model are useful for understanding nonholonomic systems generally. Wheel rotation relates to change in the position/orientation of the base. This relationship requires definition of θ_1, θ_2, X, Y, and ϕ, as well as definition of the geometric parameters R and b, as given in Figs. 1 and 2.

As will be discussed, three independent constraints, nonholonomic constraints,[2] exist among the five coordinates θ_1, θ_2, X, Y, and ϕ. Note that X and Y denote the in-plane position of the midpoint of the axle. Consider next a very small increment in the two wheel rotations, $\Delta\theta_1$ and $\Delta\theta_2$. As indicated in Fig. 2b, the no-wheel-slip assumption results in small incremental movements of the two ends of the axle of $R\Delta\theta_1$ and $R\Delta\theta_2$, respectively.

The resulting small increments become exact in the limit as the rotations become infinitesimally small:

$$\Delta\phi = \frac{R(\Delta\theta_1 - \Delta\theta_2)}{2b} \tag{1}$$

$$\Delta X = \Delta s[\cos(\phi)] \tag{2}$$

$$\Delta Y = \Delta s[\sin(\phi)] \tag{3}$$

Next, we define α and u such that $\Delta\alpha = \Delta s/R = (\Delta\theta_1 + \Delta\theta_2)/2$ and $u = (\Delta\theta_1 - \Delta\theta_2)/(\Delta\theta_1 + \Delta\theta_2)$. Considering the infinitesimal limit on these definitions, the following differential kinematic equations of motion result:

$$\frac{\mathrm{d}X}{\mathrm{d}\alpha} = R\cos\phi \tag{4}$$

$$\frac{\mathrm{d}Y}{\mathrm{d}\alpha} = R\sin\phi \tag{5}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}\alpha} = \frac{uR}{b} \tag{6}$$

Defining the three-element vector $\boldsymbol{x}$ such that $\boldsymbol{x} = [X, Y, \phi]^T$, Eqs. (4–6) may be written in the state-equation form

$$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}\alpha} = \boldsymbol{f}(\boldsymbol{x}, u) \tag{7}$$

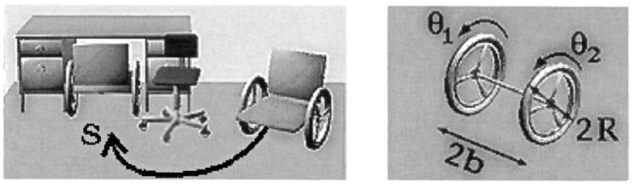

Fig. 1 Wheel rotations: θ_1 and θ_2.

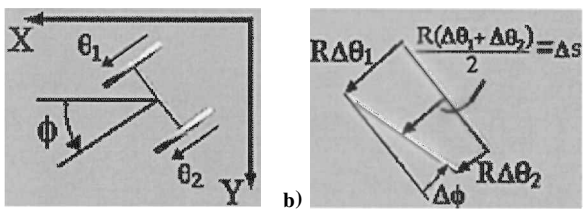

Fig. 2 In-plane position and orientation: X, Y, and ϕ.

Note that it is unnecessary to introduce the variable time t into this model. Because the nature of the nonholonomic, no-wheel-slip constraints is such that the departure of integrals of Eqs. (7) from the actual path is likely to grow with increasing α (average forward wheel rotation) rather than with time per se, this time-independent form is preferred for applications of the EKF. Specifically, dynamic (kinetic) effects on wheel slip are deemed negligible due to low operating accelerations and speeds; the order and complexity of the required model are reduced significantly due to this negligibility.

EKF

For any nonholonomic system to repeat a path effectively, it must solve the location problem; that is, it must be able to determine its current pose. This problem is made difficult because the input–output equations for a nonholonomic system are at most differential. Many early nonholonomic systems addressed this problem by simply integrating forward the differential equations of motion, a process commonly referred to as odometry or dead reckoning. In practice, this has proven very inaccurate over long paths because small errors in initial conditions, the system model, wheel slippage, and the numerical integration tend to cause estimate errors to grow as the integration progresses.

Several approaches have been taken to address this problem. Among others, Borenstein[3] and Feng and Krogh[4] have suggested various means of improving dead reckoning, including using two coordinated robots or, alternatively, an adaptive model. Such systems are quite expensive in the case of the former solution and require extensive setup in the case of the latter. More importantly, they only address portions of the dead-reckoning problem (one system addresses errors in the model, the other errors due to a nonplanar surface). Other researchers, including those from Macleod Technologies,[5] have eliminated the dead-reckoning process completely and instead use triangulation to compute the system's current pose. Such systems must always have a minimum number of beacons (reference points) available, and accuracy is limited by the accuracy of the most recent measurement.

Recently, several researchers have begun using estimation to determine the pose of the vehicle.[6–9] The present work builds on Ref. 9; it is based on the following development.

The time-independent formulation of the state equations are of the form

$$\frac{\mathrm{d}\boldsymbol{x}(\alpha)}{\mathrm{d}\alpha} = \boldsymbol{f}[\boldsymbol{x}(\alpha), u(\alpha)] + \boldsymbol{w}(\alpha) \tag{8}$$

where $\boldsymbol{x}$ is the state vector (size $n = 3$), α is the independent variable, u is the control, and $\boldsymbol{w}$ is a 3-element process-noise vector, with an associated 3×3 covariance matrix $[Q]$. Furthermore, consider that observations of the system can be made:

$$z(\alpha) = \boldsymbol{h}[\boldsymbol{x}(\alpha)] + \boldsymbol{v}(\alpha) \tag{9}$$

where $\boldsymbol{z}$ is the vector of observations (size $m = 1$) and $\boldsymbol{v}$ is an m-element measurement noise vector, with associated $m \times m$ covariance matrix $[R]$. Note that $\boldsymbol{w}$ and $\boldsymbol{v}$ are assumed to be zero-mean, Gaussian-distributed random noise processes. Additionally, they are assumed to be uncorrelated among elements and over time. Whereas it would be difficult to argue for the validity of the uncorrelated white noise assumption on $\boldsymbol{w}$ for this application, the assumption does lead to an effective algorithm because it acknowledges that there are errors in the system model and in the measurements. Moreover, experimental results given herein indicate that the distribution of actual errors, as measured in the laboratory, is consistent with predictions as derived from this model.

The estimate of the state is denoted by $E[\boldsymbol{x}(\alpha)]$, and the estimation error covariance matrix of the state is denoted by $[P(\alpha)]$ and is defined according to

$$[P(\alpha)] = E(\{\boldsymbol{x}(\alpha) - E[\boldsymbol{x}(\alpha)]\}\{\boldsymbol{x}(\alpha) - E[\boldsymbol{x}(\alpha)]\}^T) \tag{10}$$

where E represents the expectation process.[1] For discussion regarding the means by which the elements of $[P]$ are computed and used in the EKF, see Ref. 1. For a discussion of using information from more than one type of sensor in the context of the present system, see Ref. 10.

Video Observations

The present system uses charge-coupled device (CCD) cameras mounted on the wheelchair to detect passive visual cues placed at known locations throughout the environment. The horizontal locations of these cues in the image plane are the observations used by the EKF. To utilize this information, the observation equation for this measurement system must be developed.

The relationship between the position and orientation of the vehicle and the position of a cue in the image plane of the video camera mounted on the vehicle is derived by assuming that the camera can be modeled as a pin-hole camera. Using the pin-hole camera model, the image-plane location of a cue is related to physical space through the following two equations:

$$x_c = fx/z \tag{11}$$

$$y_c = fy/z \tag{12}$$

where x_c and y_c are the horizontal and vertical image-plane locations of the cue; f is the effective focal length of the camera; and x, y, and z are the location in physical space of the cue in the camera-fixed reference frame shown in Fig. 3, where O is the focal point of the camera.

Using Fig. 3b to establish some geometric relationships, we can express the horizontal image-plane location of cue P as

$$x_{cP} = f\frac{X_{P/O}\cos\phi^* + Y_{P/O}\sin\phi^*}{-X_{P/O}\sin\phi^* + Y_{P/O}\cos\phi^*} \tag{13}$$

where ϕ^* is the vehicle orientation ϕ plus a constant that relates to the fixed camera orientation on the vehicle.

Equation (13) expresses the observation equation in terms of the relationship between P and O. Introducing $\boldsymbol{C} = [C_1, C_2, C_3, C_4]^T$ to pertain to a particular one of the two onboard cameras, one of the observation equations of Eq. (9) becomes

$$x_{c_P} = C_1\frac{(X_P - X)\cos(\phi + C_4) + (Y_P - Y)\sin(\phi + C_4) + C_2}{-(X_P - X)\sin(\phi + C_4) + (Y_P - Y)\cos(\phi + C_4) + C_3} + v$$

$$\equiv g(X, Y, \phi; X_P, Y_P) + v \tag{14}$$

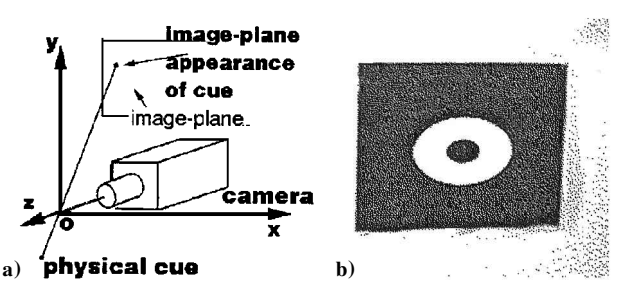

Fig. 3 Pinhole camera model and typical wall cue (located at point P).

where v is the random term as described for Eqs. (9). In practice, the parameters C_1–C_4 are determined by calibration.[9] Although X, Y, and ϕ of Eq. (14) are the states, completion of the observation equation requires X_P, Y_P, the previously measured coordinates of the detected cue. These cue locations in turn depend on cue identification, and the necessity for reliable identification motivates the development that follows.

Assigning Probability to Detected-Cue Identities

Once a cue is detected, the observation equation(s), of the form of Eq. (14), can be used, together with current estimates $E(\boldsymbol{x})$ and $[P]$ to assign a probability that the detected, located cue has a given identity. This probability assignment is based on Bayes' rule,[11] given here as

$$P(A)P(B \mid A) = P(B)P(A \mid B) \tag{15}$$

where $P(B \mid A)$ is the probability that event B occurs given that event A occurs. We consider event B to be that a particular, candidate cue, cue i (i a numerical identifier of the cue), is the one that is detected in the image. We consider event A to be that any cue indication of a particular, discernible type (in this case either white centered or black centered; see Fig. 3b) is located via image-analysis software[9] within the small interval Δx_c of horizontal camera space, which is symmetrically located about a specific point x_c in a given image. Once such a cue has been detected, it remains to determine

$$P(B \mid A) = \frac{P(B)P(A \mid B)}{P(A)} \equiv \frac{\text{NUM}}{\text{DEN}} \tag{16}$$

Consider first $P(A \mid B)$ of the numerator (NUM) of the right-hand side of Eq. (16). The probability that one particular cue, cue i, given that it is known to be detected, appears in a certain region, e.g., between x_{c1} and x_{c2} of camera space, depends on a probability density function $f_i(x_c)$ according to

$$P(x_{c2} < x_c < x_{c1}) = \int_{x_{c1}}^{x_{c2}} [f_i(x_c)]\, \mathrm{d}x_c \tag{17}$$

or, presuming $\Delta x_c = x_{c2} - x_{c1}$ is very small,

$$P(A \mid B) = f_i(x_c)\Delta x_c \tag{18}$$

Now $P(B)$ is the a priori probability that cue i is the one detected. This number is difficult to stipulate in advance; however, it is reasonable to say that there is an equal probability, call it P_{eq}, that any cue of the color/type of interest will be detected. Thus, the numerator of Eq. (16) becomes

$$\text{NUM} = P_{\text{eq}} f_i(x_c)\Delta x_c \tag{19}$$

This NUM can be interpreted as the probability that cue i is detected and that it is detected in the very small camera-space region of interest. The denominator (DEN) of Eq. (19) is the complete, a priori probability that a cue, any cue of the right color, will be detected in this same small region Δx_c, i.e., the probability of event A. Because event A could be associated with any cue, or even with a false cue detection, DEN can be written as the sum of all of the probabilities that specific cues (or a false detection) are detected in the sliver of camera-space region of interest Δx_c:

$$\text{DEN} = P(A) \approx \sum_{j=1}^{n} P_{\text{eq}} f_j(x_c)\Delta x_c + P_{\text{fd}}[f_{\text{fd}}(x_c)]\Delta x_c \tag{20}$$

where n is the number of cues deemed detectable (each of which is enumerated over j), P_{fd} is the probability of a false detection in any new sample, and $f_{\text{fd}}(x_c)$ is the probability density function of the camera-space location of the false detection given that, indeed, one does occur. Because there is assumed to be no location preference for a false detection, $f_{\text{fd}}(x_c)$ is taken to be a constant f_{fd}, defined as the reciprocal of the full range or extent of x_c.

Substituting Eqs. (19) and (20) into Eq. (16) leaves

$$P(B \mid A) = \frac{f_i(x_c)}{\sum_j f_j(x_c) + P_{\text{fd}} f_{\text{fd}} / P_{\text{eq}}} \tag{21}$$

where it is noted that $P(B \mid A)$ is the quantity of interest, i.e., the probability that cue i is the identity of the detected cue given the location of the detection. Note that $P_{\text{fd}}/P_{\text{eq}}$ of Eq. (21) is the ratio of the probability of a false detection to the probability of a particular cue's actual detection in any sample. Because most available cues are detected, P_{eq} is near unity (the image window across which search occurs continues to be expanded about the expectation location until a cue is found). If P_{fd} is taken to be approximately the probability of a detection or indication in the environment of interest when all cues are covered, it is relatively easy to establish experimentally by occluding the cues and calculating the fraction of frames where indications are found with the vehicle posed in many ways throughout the environment, about 0.05.

Probability Density Functions

The individual functions $f_j(x_c)$ are easily approximated using the estimates of X, Y, and ϕ, as described, as well as the covariance matrix $[P]$, also as described, all of which are products of the ongoing Kalman filter calculations. Requiring the function $f_j(x_c)$ to be Gaussian, we take the mean of the function $E(x_c)$ from Eq. (14) to be $g[E(X), E(Y), E(\phi); X_P, Y_P]$, where X_P, Y_P are the previously measured, physical in-plane coordinates of cue j. The variance of $f_j(x_c)$ is $E\{[x_c - E(x_c)]^2\} \equiv \sigma^2$ and is approximated by first expanding g in a Taylor series about $E(x_c)$. Retention of the linear terms only of the expansion produces

$$\sigma^2 \approx E\left\{\left[\Delta_X\left(\frac{\partial g}{\partial X}\right) + \Delta_Y\left(\frac{\partial g}{\partial Y}\right) + \Delta_\phi\left(\frac{\partial g}{\partial \phi}\right) + v\right]^2\right\} \tag{22}$$

where the partials are evaluated at $E(\boldsymbol{x})$ and at X_P, Y_P according to the in-plane physical coordinates of cue j. Also, $\Delta_X = X - E(X)$, $\Delta_Y = Y - E(Y)$, $\Delta_\phi = \phi - E(\phi)$; here $E[v^2] \equiv r^2$ is the assumed variance of the (zero-mean, uncorrelated) measurement noise. Equation (22) results in

$$\begin{aligned}\sigma^2 = {} & P_{11}\left(\frac{\partial g}{\partial X}\right)^2 + P_{22}\left(\frac{\partial g}{\partial Y}\right)^2 + P_{33}\left(\frac{\partial g}{\partial \phi}\right)^2 \\ & + (P_{12} + P_{21})\left(\frac{\partial g}{\partial X}\right)\left(\frac{\partial g}{\partial Y}\right) + (P_{13} + P_{31})\left(\frac{\partial g}{\partial X}\right)\left(\frac{\partial g}{\partial \phi}\right) \\ & + (P_{23} + P_{32})\left(\frac{\partial g}{\partial Y}\right)\left(\frac{\partial g}{\partial \phi}\right) + r^2\end{aligned} \tag{23}$$

where the quantities P_{ij} are elements of the EKF's estimate-covariance matrix $[P]$ as already defined. With the probability estimates for all possible cue identities associated with a particular indication in hand, a threshold, e.g., $P(B \mid A) = 0.95$, may be set as the condition for accepting a particular cue identity, thereby utilizing the associated sample in the estimates.

Other, more ad hoc, criteria for accepting cue identities have been considered, e.g., closest to expected camera locations, highest $P(A \mid B)$, tried over time, and rejected as inadequate. In fact, it has been the described implementation that has made apparent the need to improve on the absolute probabilities by utilizing the enhancements of the following sections. There are two reasons for the importance of this. 1) Few things are more damaging to retention of good estimates than misidentification of a cue. 2) When the estimates do degrade due to misidentification, the variances will not, by themselves, reflect this. It is, instead, the consequent decline of subsequently detected cues' identification probabilities that leads to the conclusion of severe estimate corruption.

Second Cue Attribute

The limitation, or tradeoff, associated with the described strategy for gaining cue identification may be obvious: Increasing the threshold for accepting a cue identity (clearly a desirable thing to do) risks failure to identify. Larger physical separations between placed cues helps here, but with fewer cues the frequency of observations goes down, thus reducing system position/orientation-estimate accuracy.

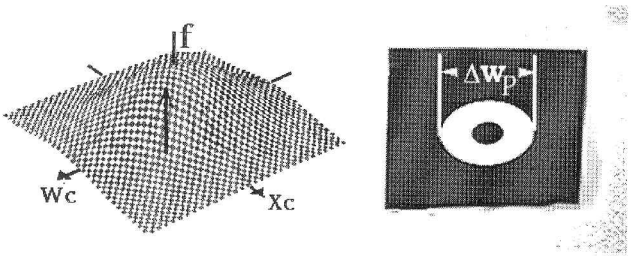

Fig. 4 Joint probability density function involving camera-space cue width and position.

Two remedies to this problem are investigated: 1) consideration of a second cue attribute, cue image-plane width, and 2) multiple cues/observations. Item 1 is considered in the present section, and item 2 is discussed subsequently.

Just as Bayes' rule was used earlier to determine the probability of a cue detected at image-plane location x_c having identity i based on the probability density function $f_i(x_c)$ together with similar functions for the other cues, $f_j(x_c)$, $j = 1, 2, 3, \ldots$, so, too, Bayes' rule can be used with combined criteria, in this instance cue width in the image plane w_c combined with x_c as used before. Figure 4 shows the joint probability density function associated with a given cue indication, using the two attributes. Representing cue j, its functional form is

$$f_j(x_c, w_c) = (2\pi)^{-1}|Q|^{-\frac{1}{2}} \exp\left[-\tfrac{1}{2}\{\boldsymbol{e}\}^T[Q]\{\boldsymbol{e}\}\right] \tag{24}$$

where $[Q]$ is the a priori covariance matrix for the two metrics, calculated from $[P]$ as will be described, and where

$$\{\boldsymbol{e}\} = [e_1, e_2]^T = \{[x_c - E(x_c)], [w_c - E(w_c)]\}^T \tag{25}$$

where $E(x_c)$ is determined from $E(\boldsymbol{x})$ as discussed earlier and $E(w_c)$ is determined from prior knowledge of the physical width of cue j, prior knowledge of the unit normal for this cue, and $E(\boldsymbol{x})$ as follows. Referring to Eq. (14), the observation equation, given by $x_{cp} = g(X, Y, \phi; X_P, Y_P) + v$, it is noted that the image-plane cue width Δx_{c_P} can be written

$$\Delta x_{c_P} \approx \left(\frac{\partial g}{\partial X_P}\right)\Delta w_X + \left(\frac{\partial g}{\partial Y_P}\right)\Delta w_Y \tag{26}$$

where Δw_X and Δw_Y are components of the physical width of the cue as projected onto the X and Y physical (room) axes, respectively. They are determined from the (input) total physical width of the cue Δw_P combined with the (also input) unit normal to the cue, $\boldsymbol{e}_P \equiv e_X\boldsymbol{i} + e_Y\boldsymbol{j}$ ($\boldsymbol{i}$ and $\boldsymbol{j}$ aligned with the X and Y axes), according to

$$\Delta w_X = -\Delta w_P e_Y \tag{27}$$

$$\Delta w_Y = \Delta w_P e_X \tag{28}$$

Thus, referring to Eq. (26), we may write

$$\Delta x_{c_P} = h(X, Y, \phi; X_P, Y_P, \Delta w_P, e_X, e_Y) + \varepsilon \tag{29}$$

where ε is a random term (similar to v), which is taken here to be negligibly small. Defining $[Q]$ as $E(\{\boldsymbol{e}\}\{\boldsymbol{e}\}^T)$, it is straightforward to derive its four elements according to Eq. (25), in terms of $E(X)$, $E(Y)$, $E(\phi)$ as well as $[P]$ by noting

$$e_1 = \Delta_X\left(\frac{\partial g}{\partial X}\right) + \Delta_Y\left(\frac{\partial g}{\partial Y}\right) + \Delta_\phi\left(\frac{\partial g}{\partial \phi}\right) + v \tag{30}$$

$$e_2 = \Delta_X\left(\frac{\partial h}{\partial X}\right) + \Delta_Y\left(\frac{\partial h}{\partial Y}\right) + \Delta_\phi\left(\frac{\partial h}{\partial \phi}\right) \tag{31}$$

where, as before, all partials are evaluated at $E(\boldsymbol{x})$. Finally, Bayes' rule is applied in a direct extension of Eq. (21) with

$$P(B \mid A) = \frac{f_i(x_c, w_c)}{\sum_j f_j(x_c, w_c)} + \frac{P_{\text{fd}} f_{\text{fd}}(x_c, w_c)}{P_{\text{eq}}} \tag{32}$$

where the function $f_{\text{fd}}(x_c, w_c)$ is taken to be constant across the range of (x_c, w_c) space deemed possible for false detections.

Experimental Results

The apparatus used to produce experimental results is shown in Fig. 5. In addition to being able to read the ongoing encoder angles for each of the two rear drive wheels, the apparatus has two cameras mounted beneath the seat, which are able to detect cues on walls.

One kind of test performed using the apparatus, mentioned earlier, is the extent of agreement between actual, measured vehicle position error and the distribution of error expected from EKF-based probability density functions. To test this, the chair was stopped at approximately 50 locations at various junctures during several different repeat trajectories, and its position error normal to the vehicle direction was measured (using a plumb bob dangling down to the floor) against the counterpart taught position. The system's estimate of the discrepancy between these two position components was subtracted from the observed quantity. This figure, taken to be the error in the vehicle's position estimate, was then normalized by the standard deviation computed using the EKF for this error component. The ideal result of a unity-variance Gaussian distribution is shown plotted against the actual distribution of errors in Fig. 6. This result adds to the confidence of basing cue-identity probabilities on the EKF's error covariances.

The most straightforward means by which to report the advantage gained by introducing the second attribute, cue width, into consideration is to report on the number of rejected measurements observed during a pair of otherwise identical runs, one run using x_c only and the second using both criteria x_c and w_c. For both runs, the requirement for acceptance of a cue identity would be a probability greater than 0.95.

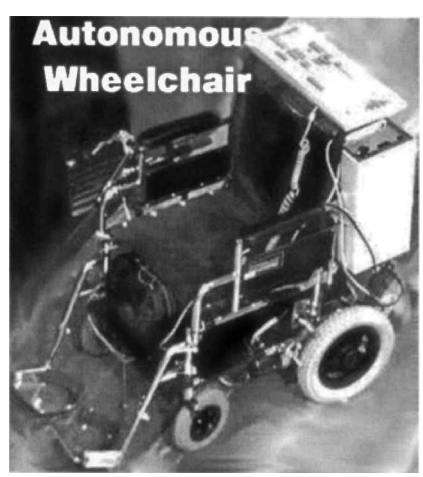

Fig. 5 Experimental apparatus.

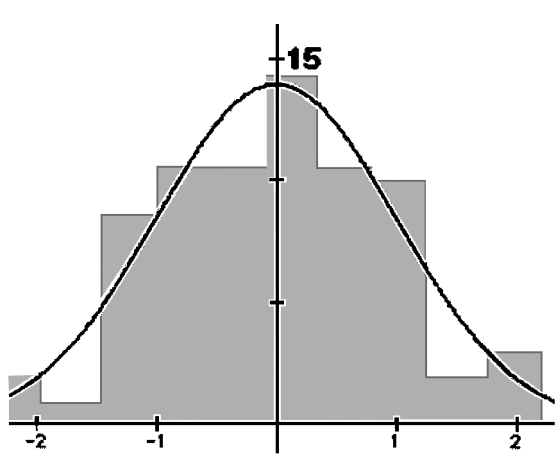

Fig. 6 Histogram of normalized errors.

In 15 separate repeat runs (long trajectories, with path lengths exceeding 30 ft, previously taught by the human operator), on average, the use of the dual criterion reduced the incidence of rejected cue information by about 80% compared with the same run using position only. Where the single attribute is used, the average probability associated with the correct cue identity is about 0.939; where both attributes are used, this figure goes to 0.986.

Moreover, whereas instances of wrong cue-identity assignments were reasonably common prior to implementation of this measure, following implementation not a single case of false identification has occurred with the same cue spacing. For more information on this and further experimental error distributions, see Ref. 12.

Using Multiple Cues

Although some considerable advantage is realized by involving the second attribute, for purposes of safety and reliability there is still a desire to at once significantly increase the acceptance-probability level to perhaps 0.995 while doubling or tripling the density of cues on the wall. Because cue position and width seem to be the only reasonably reliable attributes that are chair-location dependent, we turn to a somewhat distinct strategy: multiple cues.

It is possible, given adequately dense positioning of cues on the wall, reliably to acquire two or more indications in any given image. Such an acquisition permits the multiple-cue-identity probabilities to be assigned based on the estimate covariance matrix $[P]$ of the EKF, analogous to the single-cue probabilities already discussed.

When N cues are detected in a given image, the counterpart to $f_j(x_c)$ of Eq. (21) or $f_j(x_c, w_c)$ of Eq. (32) is denoted here by $f_J(x_c^1, x_c^2, \ldots, x_c^N)$, $J = 1, 2, 3, \ldots, M$. Here, $f_J(x_c^1, x_c^2, \ldots, x_c^N)$ represents the value of the joint probability density function associated with the Jth of M possible sets of physical identities of the N cues detected, including the possibility of up to N false detections. Clearly, several criteria will limit the number M of possible cue identities. These include cue type, left-to-right order, and so on.

The counterpart to the Bayes rule results of Eqs. (21) and (32) becomes, for this case,

$$P(B \mid A) = \frac{f_I\left(x_c^1, x_c^2, \ldots, x_c^N\right)}{\sum_J f_J\left(x_c^1, x_c^2, \ldots, x_c^N\right)} \tag{33}$$

This is the probability of identity set I, given the N observed camera-space cue locations. The actual calculation of f_J makes use of the fact that the individual events, $x_c^1, x_c^2, \ldots, x_c^N$, are not independent. Thus, f_J is found according to

$$f_J\left(x_c^1, x_c^2, \ldots, x_c^N\right) = f_1\left(x_c^1\right) f_2\left(x_c^2 \mid x_c^1\right), \ldots, f_N\left(x_c^N \mid x_c^1, x_c^2, \ldots, x_c^{N-1}\right) \tag{34}$$

where the indices 1, 2, 3, ..., refer to the previously ordered, hypothesized cue-identity sets and associated probability density functions connected with the Jth considered sequence. Thus, $f_2(x_c^2 \mid x_c^1)$, for example, is the probability density function associated with the second sequenced cue (for sequence J) given that the estimates $E(X)$, $E(Y)$, $E(\phi)$, and $[P]$ are first updated by using the sample x_c^1 with the assumption that this sample corresponds with the first cue identity in sequence J.

If the observation equation [Eq. (14)] were linear in X, Y, and ϕ, then the function f_J of Eq. (34) would be unaffected by the order in which the observations identified with J are arranged. However, with the nonlinearity there is some small discrepancy depending on this order, indicating the approximate nature of the probability of Eq. (33).

As an indication of the effectiveness of this new approach, consider the following. When $N = 2$, in a series of recent typical trials, the average calculated probability associated with the correct interpretation of the two cues is about 0.9999, whereas, using the single-cue position only criterion, the average probability during a typical run associated with the correct cue identity is about 0.94. Using the single-cue dual criteria, this rises to 0.986. Clearly, it will be worthwhile extending the preceding development to multiple cues entailing the dual criteria.

Summary

Just as the combination of location estimation and attribute identification is considered in certain missile-guidance applications,[13] the paper presents an example of the same combination of considerations with a potentially every day, domestic application, the automatic guidance of a power wheelchair for severely disabled individuals.[14] Illustrated is a series of increasingly effective ways in which a byproduct of the EKF estimation algorithm, the error covariance matrix, can be used to identify with high certainty the particular cues detected via video to upgrade position/orientation estimates.

In practice, the single most common reason for failure with the experimental apparatus has been insufficient cue density, misidentified cues, and/or inadequate probability to disambiguate cue identity. Such failure has been associated with both inadequate position/orientation accuracy of the estimates due to too few observations per unit distance traveled; it also is associated with false cue identities. Thus, the combination of single-attribute, single-cue identification and more ad hoc methods of identification (such as the closest-to-prediction criterion) is what has motivated this work. The promise of using multiple indications in a single image, is illustrated, as well as multiple cue attributes, to allow both denser cue representation in the environment and significant improvement in the probability of correctly identifying detected cues. With implementation of these measures system reliability is approaching 100%. (Note that, in practice, initialization could be problematic unless there is some assurance of accurate initial positioning; hence, we have devised special-purpose, distinct cue forms,[15] which provide good initialization if they are in view.)

Conclusion

The work carries important implications not only for the present system, with its requisite positioning of artificial marks on the wall, but it may be the key to subsequent efforts, which make use of the easily detected vertical edges that are abundant naturally in the fabricated settings for which this class of vehicles (as well as others such as floor-maintenance vehicles) are likely to operate. (Digital video and discussion of other aspects of the system discussed herein is available on the World-Wide Web, URL: http://www.nd.edu/NDInfo/Research/sskaar/Home.html.)

Acknowledgment

This work is supported, in part, by the Office of Naval Research, U.S. Department of the Navy, N00014-95-1-G010.

References

[1] Gelb, A. (ed.), *Applied Optimal Estimation*, MIT Press, Cambridge, MA, 1974, pp. 182–203.

[2] Meirovitch, L., *Methods of Analytical Dynamics*, McGraw–Hill, New York, 1970.

[3] Borenstein, J., "The Clapper: A Dual-Drive Mobile Robot with Internal Correction of Dead-Reckoned Errors," *Proceedings of the IEEE International Conference on Robotics and Automation*, Inst. of Electrical and Electronics Engineers, New York, 1994, pp. 3085–3090.

[4] Feng, D., and Krogh, B. H., "Dynamic Steering Control of Conventionally-Steered Mobile Robots," *Proceedings of the IEEE International Conference on Robotics and Automation*, Inst. of Electrical and Electronics Engineers, New York, 1990, pp. 390–395.

[5] "CONAC: Computerized Opto-Electonic Navigation and Control," Macleod Technologies, Inc., Chelmsford, MA, 1994.

[6] Kleeman, L., "Optimal Estimation of Position and Heading for Mobile Robots Using Ultrasonic Beacons and Dead-Reckoning," *Proceedings of the 1992 IEEE Conference on Robotics and Automation*, Inst. of Electrical and Electronics Engineers, New York, 1992, pp. 2582–2587.

[7] Leonard, J. J., and Durrant-Whyte, H. F., *Directed Sonar Sensing for Mobile Robot Navigation*, Kluwer Academic, Dordrecht, The Netherlands, 1992.

[8] Crowley, J. L., and Chenavier, F., "Position Estimation for a Mobile Robot Using Vision and Odometry," *Proceedings of the 1992 IEEE Conference on Robotics and Automation*, Inst. of Electrical and Electronics Engineers, New York, 1992, pp. 2588–2593.

[9] Baumgartner, E. T., and Skaar, S. B., "An Autonomous, Vision-Based Mobile Robot," *IEEE Transactions on Automatic Control*, Vol. 39, No. 3, 1994, pp. 493–502.

[10] Baumgartner, E. T., and Yoder, J. D., "Enhancing Mobile Robot Location Estimates Using Sensor Fusion," *Proceedings of the ASME Dynamic Systems and Control Division*, DSC-Vol. 57-1, American Society of Mechanical Engineers, New York, 1995, pp. 137–144.

[11] Walpole, R. E., and Myers, R. H., *Probability and Statistics for Engineers and Scientists*, 2nd ed., Macmillan, New York, 1978.

[12] Arriola, H., "Investigations of Methods for Cue Identification Utilized by an Autonomous Wheelchair," M.S. Thesis, Dept. of Aerospace and Mechanical Engineering, Univ. of Notre Dame, Notre Dame, IN, Jan. 1997.

[13] Junkins, J. L., *An Introduction to Optimal Estimation of Dynamical Systems*, Sijthoff and Noordhoff, Alphen Aan Den Rijn, The Netherlands, 1978.

[14] Yoder, J. D., Baumgartner, E. T., and Skaar, S. B., "Initial Results in the Development of a Guidance System for a Powered Wheelchair," *IEEE Transactions on Rehabilitation Engineering*, Vol. 4, No. 3, 1996, pp. 143–151.

[15] Yoder, J. D., "Advanced Topics for the Navigation of an Automatically Guided Wheelchair System," Ph.D. Dissertation, Dept. of Aerospace and Mechanical Engineering, Univ. of Notre Dame, Notre Dame, IN, April 1996.