Sensor Action Planning driven by Uncertainty. Application to Object Location with Robust Local Sensors in a Nuclear Environment.

J. De Geeter[1], J. De Schutter[2], M. Decréton[1], H. Van Brussel[2]

[1]SCK•CEN Belgian Research Centre for Nuclear Energy,
Teleoperation Project, Boeretang 200, B-2400 Mol, Belgium.

[2]K.U.Leuven, Dept. of Mechanical Engineering, Div. PMA,
Celestijnenlaan 300B, 3001 Heverlee, Belgium
email : jan.degeeter@sckcen.be

Abstract

This paper presents an integrated Bayesian solution to the problem of object location estimation, object recognition and sensor action planning under uncertainty. The emphasis is on finding the best next sensing action.

The method uses elementary notions from Bayesian decision theory. The best action is found as the one that optimises the expected value of a utility function, which is the logarithm of the volume of the uncertainty ellipsoid around the estimate of a target position.

An example shows that this method is capable of controlling the sensing actions of an ultrasonic sensor mounted on a robot, where the target is to accurately position a drill on a cylinder before drilling a hole.

The presented algorithm is easy to apply and computationally tractable.

Keywords: sensor planning, uncertainty, Bayesian decision theory, ultrasonic sensing, teleoperation.

1 Introduction

Finding a good viewpoint is a problem of everyday life. If you are uncertain as to where you are in a large forest, you will purposefully look for information. However, deciding where to look and how much effort to spend to get more certainty on your situation is a nontrivial problem: an action which you think will yield a lot of information may be very expensive, and drive you further away from your goal, whereas less informative actions may divert you a lot less from what you assume is the right track ...

This paper solves a similar problem in an industrial environment: given our current incomplete and uncertain knowledge on the identity and location of objects, and given a set of noisy sensors, the problem is to develop an intelligent sensing behaviour that reduces the uncertainty on the goal as fast as possible. If the task is e.g. to drill a hole in a workpiece with a well-defined tolerance, then the goal is the drill position on the workpiece.

International Journal of Computing Anticipatory Systems, Volume 2, 1998
Ed. by D. M. Dubois, Publ. by CHAOS, Liège, Belgium. ISSN 1373-5411 ISBN 2-9600179-2-7

Figure 1: *A nuclear work cell.* Operations in such a cell are very similar to what happens in a ordinary factory workshop, except that work cannot be done hands-on due to radiation. The operator is looking through a lead glass window into the cell. The cell shown here is equipped with traditional mechanical master-slave manipulators.

We propose an entirely Bayesian decision-theoretic solution to this problem, which finds, at each time step, the next action that minimises the expected value of a utility function. This utility function relies on a measure of uncertainty (or conversely, a measure of information), equal to the volume of the uncertainty ellipsoid of the covariance matrix. The sensing behaviour can be influenced by adding extra terms to this utility function, e.g. expressing risk, elapsed time, etc.

This solution establishes sensing strategies which people usually agree upon as being acceptable. We do not claim to have *the unique optimal* strategy, which does not exist. However, we do claim that we have a simple strategy which is computationally tractable, and which can be integrated easily with Bayesian algorithms for object location and recognition, as will be shown.

This paper first puts the active sensing problem in the context of our application, Section 2. Section 3 gives a flavour of the research being done in the field of sensor viewpoint planning. Section 4 then outlines our approach to sensor planning, and briefly reviews our previous work on Bayesian object location and recognition, necessary for understanding the remainder of this paper. Section 5 presents our decision theoretic approach to solving the viewpoint planning problem. Section 6 illustrates this with some examples.

2 Towards a Decision-Support for the Location and Recognition of Objects in a Nuclear Environment

The goal of this work is to automate some parts of teleoperated tasks for which accurate geometric information is important, such as assembly, cutting, grasping or drilling. This work concentrates on applications in high radiation fields in which human intervention is impossible, conditions which frequently occur in nuclear work cells, as shown Fig. 1, or in future fusion reactors.

Under these extreme radiation conditions, the choice of sensors becomes very critical (Decréton 1995): The most robust sensors are those with simple transducers and remote electronics, such as optical detectors, some ultrasonic sensors, some types of force sensors

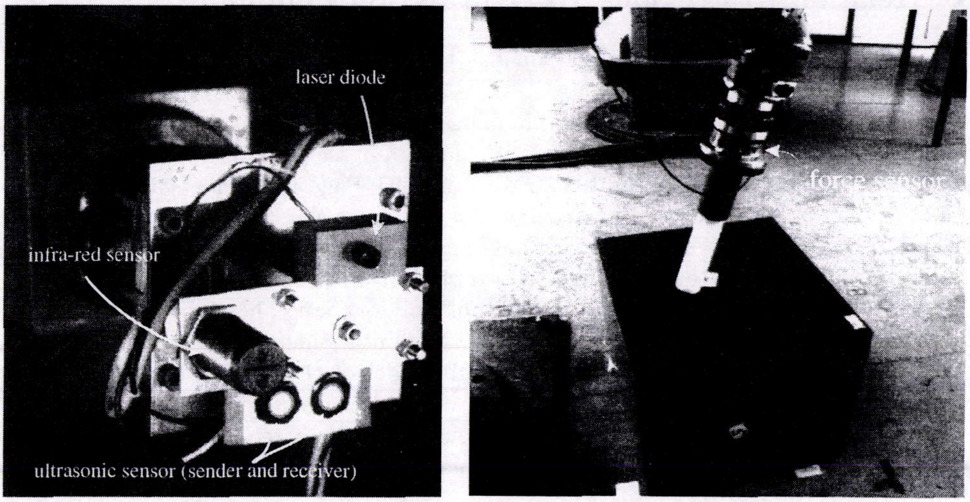

Figure 2: *Robust radiation-resistant sensors.* The left figure shows a robot equipped with an ultrasonic sensor and an infra-red optical detector. The ultrasonic sensor measures the distance to a surface. The infra-red sensor returns an output voltage which is a function of the intensity of the reflected light, and is used to detect edges. The white probe (right figure) is used to explore the environment by touch, very much like the visually impaired. A force sensor measures the forces on this probe in six dimensions (three forces, three moments)

etc., Fig. 2. In contrast to this, the performance of more complex sensors such as CCD cameras, degrades quickly.

These robust sensors however return only *local* measurements, such as the distance to a point on a surface, in contrast to cameras which return a huge amount of information in one shot, of which usually however only a small amount is relevant. Image processing requires immense computing power mainly to throw away irrelevant data. In contrast to this, the data of local sensors is necessarily more sparse and very precious. Therefore, our approach is exactly the opposite: to purposefully gather information so as to maximise the information content of each measurement. This approach is often referred to as *active sensing* (Bajcsy 1988).

Local data is, for a human operator, far more difficult to interpret than the data of more anthropomorphic sensors such as cameras. Therefore, it is our aim to construct a *decision aid for the human operator*, that assists him in estimating the position of objects from the sensor data, in recognising objects, and in planning the next measurement. The methods proposed in this paper are the core of such an operator support.

Note that, although our work is driven by nuclear applications, obviously the developments are also useful to every other application domain where sensor data is not available for free. In addition, local sensors can also be useful in non-hazardous domains due to space or weight limitations: these small sensors can be added to every tool, or even the tool itself can be effectively used a a measurement probe. This opens new perspectives for sensor-based robotics, as one of the major criticisms up to now is the lack of robustness of sophisticated sensors.

3 Relation to Previous Work in Action Planning

The problem of finding the best action has been solved in many ways and for many purposes.

Many have studied the problem of planning the optimal trajectory for a mobile robot, *from start to finish*, in order to reach a certain goal state, see e.g. (Barraquand and Ferbach 1995; Bouilly, Siméon, and Alami 1995; Erdmann 1995; Page and Sanderson 1995). Others prefer to plan only one action or *one time step ahead*, see e.g. (Das, Beni, and Hackwood 1992; Kristensen 1995). This approach only maximises the reward of the next action, at the expense of the risk that the total reward of all future actions may be lower than for another action. Despite this, we prefer to look only one step ahead for various reasons. First, optimisation from start to finish is often computationally intractable (Barraquand and Ferbach 1995). Second, a fairly accurate model of the environment is needed to compute a useful action plan. Unfortunately, our task is exactly to model (part of) this environment. In this case, the action plan should be recalculated each time the model is updated. In addition, computing the action plan beforehand is not useful if the world is dynamic.

Uncertainty can be described with *error bounds*, see e.g. (Erdmann 1995; Page and Sanderson 1995) or with *probability density functions* (PDF), see e.g. (Barraquand and Ferbach 1995; Borghi and Caglioti 1995; Das, Beni, and Hackwood 1992; Whaite and Ferrie 1997; Kristensen 1995). In both cases, finding the best strategy amounts to playing a game, where the opponent is Nature. If Nature's behaviour can be described with a PDF, finding the optimal strategy reduces to an optimal control problem (LaValle 1995). If uncertainty is described by sets, calculations become more difficult as the dimension of the problem increases, see (Hager, Engelson, and Atiya 1993) for a discussion. In addition, stochastic methods usually give better point estimates, while the set based estimation methods are better at calculating bounds on the solution.

There can be uncertainty on the state of the world, on the dynamics of the robot system and on the measurement of the sensing system. (Bouilly, Siméon, and Alami 1995) assume that the world is known exactly. Many authors implicitly assume a quite accurate world model, as discussed above. The presented solution can handle the most general case.

Most of the solutions in literature only deal with situations in which there are a discrete number of world states and a discrete number of actions, see e.g. (Kristensen 1995). Our approach can handle discrete sets of continuous states, and a discrete set of continuous actions.

The presented approach to action planning is most closely related to the methods proposed by (Borghi and Caglioti 1995; Whaite and Ferrie 1997; Kristensen 1995). These solutions (i) can cope with uncertainty in the environment, the sensor and the dynamics of the positioning device, (ii) look only one step ahead. As in (Borghi and Caglioti 1995; Whaite and Ferrie 1997; Subrahmonia, Cooper, and Keren 1996), we use the determinant of the covariance matrix as a measure for the information in our utility function. At present we do not yet consider other costs in the utility function. See e.g. (Cook, Gmytrasiewicz, and Holder 1996) for an application with a more complex utility function.

4 Our Approach to Object Location and Recognition

This section briefly reviews the solution we adopted for locating and recognising objects. For more details, the reader is referred to (De Geeter et al. 1996; De Geeter et al. 1997; De Geeter et al. 1997).

Sensor measurements contain two sources of uncertainty: (i) uncertainty on the measurement value (usually called sensor noise) and (ii) uncertainty or ambiguity on the origin of the measurement. i.e. uncertainty as to from which feature from which object the measurement originates. The Kalman filter (KF), Section 4.1 deals with the sensor noise, while the multiple hypothesis tree (MHT), Section 4.2 deals with the ambiguity.

4.1 Object Location: the Smoothly Constrained Kalman Filter

This section only briefly reviews the Kalman filter (KF), to fix notations. For more details on the KF, the reader is referred to (Bar-Shalom and Li 1993; Gelb 1974; Kalman 1960; Sorenson 1985).

The task of the estimator is to calculate the position of an object together with its variance, given the measurement, and given a hypothesis on the origin of the measurement. A KF is a linear stochastic weighted recursive least squares estimator, which is perfectly suited for this task if all equations are linear. Unfortunately, nonlinearities can prevent the estimate from converging to the true value. In our application, objects are described as a collection of simple geometrical features, such as planes or cylindrical surfaces, with constraints defining their relative position. The state of an object is uniquely described by a vector $\boldsymbol{x}$, called the state vector. A constraint is any relation that exists between state variables of $\boldsymbol{x}$, such as the distance between the two planes of a cylinder. The SCKF is an evolution of the classic KF, able to integrate nonlinear constraints into the estimate, while avoiding to corrupt the estimate with an important bias due to linearisation errors (De Geeter et al. 1997; De Geeter et al. 1996).

Input to the KF on time step i is (i) the KF estimate $\hat{\boldsymbol{x}}_{i-1}$ calculated on the previous time step, together with its covariance matrix $\widehat{\boldsymbol{P}}_{i-1}$, and (ii) the new measurement $\boldsymbol{z}_i$ with the *measurement equation*

$$\boldsymbol{z} = \boldsymbol{H}_i \boldsymbol{x} + \boldsymbol{\rho}_i, \qquad (1)$$

where $\boldsymbol{H}_i$ is the *measurement matrix*, and $\boldsymbol{\rho}_i$ is zero-mean normally distributed white noise $N(0, \boldsymbol{R})$.

$\boldsymbol{H}_i$ and $\boldsymbol{\rho}_i$ are defined by the chosen sensing action a_i. this dependency is not explicitly shown in oder not to overload notations. Is is however through this dependency that the sensing planning strategy can influence the result of the estimator.

Output from the KF is (i) the new estimate $\hat{\boldsymbol{x}}_i$, $\widehat{\boldsymbol{P}}_i$, (ii) the innovation $\boldsymbol{\nu}_i$ with variance $\boldsymbol{S}_i$, calculated as follows:

$$\begin{aligned}
\boldsymbol{\nu}_i &= \boldsymbol{z}_i - \boldsymbol{H}_i \hat{\boldsymbol{x}}_{i-1} , && (2)\\
\boldsymbol{S}_i &= \boldsymbol{R}_i + \boldsymbol{H}_i \widehat{\boldsymbol{P}}_{i-1} \boldsymbol{H}_i^T , && (3)\\
\boldsymbol{K}_i &= \widehat{\boldsymbol{P}}_{i-1} \boldsymbol{H}_i^T \boldsymbol{S}_i^{-1} , && (4)\\
\hat{\boldsymbol{x}}_i &= \hat{\boldsymbol{x}}_{i-1} + \boldsymbol{K}_i \boldsymbol{\nu}_i , && (5)\\
\widehat{\boldsymbol{P}}_i &= (\boldsymbol{I} - \boldsymbol{K}_i \boldsymbol{H}_i) \widehat{\boldsymbol{P}}_{i-1} , && (6)
\end{aligned}$$

where $\tilde{\boldsymbol{K}}_i$ is the *Kalman gain* matrix.

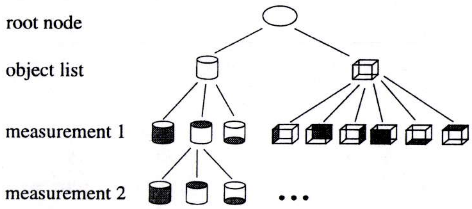

Figure 3: *Definition of the MHT.* The two top levels are respectively a level with only one root node, and a level containing a list of object models relevant to the current task, here only a cylinder and a box. Every node on one (regular) level corresponds to a hypothesis about the origin of a particular measurement (i.e. the object feature corresponding to the measurement). This feature is shaded in grey.

4.2 Object Recognition: The Bayesian multiple hypothesis tree

The multiple hypothesis tree (MHT) keeps track of different local scene models, each corresponding to a set of hypotheses on measurement-feature associations, until there is sufficient evidence to discard some of them.

The MHT is defined as follows, Fig. 3. The top level consists only of a (dummy) root node. Immediately below is an exhaustive list of all object models relevant to the current task, in this case only a cylinder and a box. Only one of these models corresponds to the object present in the real scene. Each regular *level* in the MHT corresponds to one measurement. Each *node* on a level corresponds to a hypothesis about the origin of this measurement (i.e. the object feature corresponding to the measurement). This feature is shaded in grey in Fig. 3. Each node contains an estimate of the location of the assumed object. The set of *leaf nodes* corresponds to the set of hypotheses currently under consideration. A *path* from a particular leaf node to the top node of the tree describes the set of hypotheses underlying the estimate of the leaf node, i.e. the history of an estimate, going back in time as one goes up in the tree.

This paper concentrates on problems in which the object to locate is known to be one out of a list of object models, so each level in the MHT contains exactly one correct hypothesis.

This tree is pruned by checking the residual error of the estimate corresponding to each leaf node. This residual error is equal to the Summed Normalised Innovation Squared (SNIS) (De Geeter et al. 1997), defined as follows:

$$\sigma_i^{\mathrm{SNIS}} = \sum_{k=1}^{i} \boldsymbol{\nu}_k^T \boldsymbol{S}_k^{-1} \boldsymbol{\nu}_k \ , \tag{7}$$

which is χ^2-distributed with $\sum_{k=1}^{i} l_k$ degrees of freedom (dof), where l_k is the number of statistically independent measurements in vector $\boldsymbol{z}_k$. Note that σ_i^{SNIS} can be calculated from $\sigma_{i-1}^{\mathrm{SNIS}}$, and $\boldsymbol{\nu}_i$ and $\boldsymbol{S}_i^{-1}$ which are calculated by the KF.

The probability p_j of each leaf node j being the correct one is simply proportional to the probability of the value of σ_j^{SNIS} according to its χ^2-distribution, with a normalisation to make sure that the probabilities on one level of the tree sum to 1, as the list is exhaustive.

	s_1	s_2	$\dots$	s_n
a_1	u_{11}	u_{12}	$\dots$	u_{1n}
a_2	u_{21}	u_{22}	$\dots$	u_{2n}
$\vdots$				
a_m	u_{m1}	u_{m2}	$\dots$	u_{mn}

Table 1: The *decision table* contains the utility u_{ij} of an action a_i if the world is in state s_j.

5 Decision Making Under Uncertainty

The MHT summarises at each time step the current knowledge on the state of the world, and hence serves as the basis for planning sensing actions.

Section 5.1 first explains the decision-theoretic approach for the simple classic case where a choice has to be made from a discrete set of actions, and the world can only be in a discrete set of states. Section 5.2 then extends this method to the case corresponding to the lowest level of the MHT: the choice from a discrete set of continuously parametrised actions, where the state of the world can be any of a discrete set of continuously parametrised states.

5.1 Finite Set of Actions, Finite Set of States

A choice must be made from a set of actions $\mathcal{A} = \{a_1, a_2, \dots a_m\}$, but the desirability of each of these actions depends on "the state of nature", which is a value from the set $\mathcal{S} = \{s_1, s_2, \dots, s_n\}$. The states in $\mathcal{S}$ are mutually exclusive and the set $\mathcal{S}$ is exhaustive for the current decision making problem. Here we assume that the decision maker has available a probability distribution $\mathcal{P} = \{p_1, p_2, \dots, p_n\}$ over the set $\mathcal{S}$, describing the probability of each possible state of nature actually being the true one.

Each action-state pair (a_i, s_j) has a consequence or outcome. We assume here that the preference of the decision maker for each of these outcomes is expressed by a utility function. Then, the decision problem can be summarised as in table 1. For action a_j, the expected utility $ua_j = p_1 u_{j1} + p_2 u_{j2} + \dots + p_n u_{jn}$.

We define as the "best" action the one that maximises the expected utility, or $a^* = \max_{a_j} ua_j$. Other decision criteria than the maximum utility exist, see (Duncan Luce and Raiffa 1957) for a discussion.

5.2 Finite Set of Continuously Parametrised Actions, Finite Set of Continuously Parametrised States

In this case, a choice must be made from a discrete set of continuously parametrised actions $\mathcal{A} = \{a_1(\boldsymbol{v}_1), a_2(\boldsymbol{v}_2), \dots a_m(\boldsymbol{v}_m)\}$, where $\boldsymbol{v}_j$ is a vector of continuously varying parameters. The type and number of these parameters may differ over the actions. The state is one out of a finite set of continuously parametrised states $\mathcal{S} = \{s_1(\boldsymbol{x}_1), s_2(\boldsymbol{x}_2), \dots, s_n(\boldsymbol{x}_n)\}$, where $\boldsymbol{x}_j$ is a vector of continuously varying parameters. The type and number of these parameters may differ over the states. Again, we assume that a probability distribution $\mathcal{P}$ is available over the set $\mathcal{S}$, $\mathcal{P} = \{p_1, p_2, \dots, p_n\}$. In addition, we assume that a PDF $p_{xj}(\boldsymbol{x}_j)$ over the parameter $\boldsymbol{x}_j$ of each state $s_j(\boldsymbol{x}_j)$ is available.

Again, several definitions of "best" action are possible. We consider as best action $a^*(\boldsymbol{v}^*) = \max_{a_j} \left(\max_{\boldsymbol{v}_j} ua_j\right)$, i.e. first, for each type of action a_j, the parameters $\boldsymbol{v}_j^*$

are determined which maximise the expected utility, and second, this action $a^*(\boldsymbol{v}^*)$ from $a_j(\boldsymbol{v}_j^*)$ is selected which has the maximum expected utility.

5.3 Utility Function

Our objective here is to reduce as fast as possible the uncertainty on a goal location, e.g. the position where to drill the hole. In the most general case, this goal position $\boldsymbol{x}^g$ is not identical to the estimated state $\boldsymbol{x}$, but is calculated from it with a transformation $\boldsymbol{t}$: $\boldsymbol{x}^g = \boldsymbol{t}(\boldsymbol{x})$. The dimension of vector $\boldsymbol{x}^g$ is smaller then or equal to the dimension of $\boldsymbol{x}$ (or $\boldsymbol{x}$ is usually a non-minimal state description of the goal position[1]).

A measure for the uncertainty on the goal location is the volume of the 1-σ uncertainty ellipsoid of the PDF. The uncertainty ellipsoid is defined by the covariance matrix of the PDF of the goal location as follows.
The estimated goal location $\hat{\boldsymbol{x}}_i^g$ and its covariance are calculated from the state estimate:

$$\begin{aligned} \hat{\boldsymbol{x}}_i^g &= \boldsymbol{t}(\hat{\boldsymbol{x}}_i) \, , &(8)\\ \widehat{\boldsymbol{P}}_i^g &= \boldsymbol{t}_i \widehat{\boldsymbol{P}}_i \boldsymbol{t}_i^T \, , &(9) \end{aligned}$$

where $\boldsymbol{t}_i = \left.\frac{\partial \boldsymbol{t}}{\partial \boldsymbol{x}}\right|_{\hat{\boldsymbol{x}}_i}$ is the Jacobian of transformation $\boldsymbol{t}$. The singular value decomposition of $\widehat{\boldsymbol{P}}_i^g$ is given by

$$\widehat{\boldsymbol{P}}_i^g = \boldsymbol{U}_i \boldsymbol{D}_i \boldsymbol{U}_i^T \, , \qquad (10)$$

where $\boldsymbol{D}_i$ is a diagonal matrix of the singular values, and $\boldsymbol{U}_i$ is a rotation matrix. The square root of a singular value corresponds to the length of a principal axis, while the orientation of these principal axes is given by $\boldsymbol{U}_i$. The volume of the ellipsoid is given by $\det(\widehat{\boldsymbol{P}}_i^g)$, which is also equal to the product of the singular values of $\widehat{\boldsymbol{P}}_i^g$.

We then define the utility of an action ua_j as the logarithm of the expected value of the volume of the uncertainty ellipsoid on the next time step:

$$ua_j = \log\det(\widehat{\boldsymbol{P}}_{i+1}^g) \, , \qquad (11)$$

which is a function of the action: expanding expression (11) with Eqs. (9,6,4), yields:

$$\begin{aligned} ua_j &= \log\det(\boldsymbol{t}_i \widehat{\boldsymbol{P}}_{i+1} \boldsymbol{t}_i^T) \, , &(12)\\ &= \log\det(\boldsymbol{t}_i(\boldsymbol{I} - \widehat{\boldsymbol{P}}_i \boldsymbol{H}_i^T \left(\boldsymbol{R}_i - \boldsymbol{H}_i \widehat{\boldsymbol{P}}_i \boldsymbol{H}_i^T\right)^{-1} \boldsymbol{H}_i)\widehat{\boldsymbol{P}}_i \boldsymbol{t}_i^T) \, . &(13) \end{aligned}$$

Now, the Jacobian $\boldsymbol{H}_i$ depends on (i) the type of action, (ii) the parameters of this action, (iii) the state hypothesis. This dependency is exploited here. Note that, if the measurement equation is linear, $\boldsymbol{H}_i$ is independent of the state estimate, and hence the complete sensing strategy can be planned beforehand!

5.4 Discussion

This type of utility function is well-known in literature on optimal experiment design for system identification, see e.g. (Goodwin and Payne 1977; Mehra 1974; Zarrop 1979). There, the utility or cost criterion is usually written as $\log\det \boldsymbol{M}^{-1}$, where $\boldsymbol{M}$ is the

[1]Note that the description of an object as a collection of features is already a non-minimal state description for the object, as explained in Section 4.1, and illustrated in Section 6.1

Fisher Information Matrix, which is the inverse of the covariance matrix for Gaussian distributions. This cost criterion leads to the so-called D-optimality. Other frequently used cost criteria based on the Fisher Information Matrix are trace($\boldsymbol{W}\boldsymbol{M}^{-1}$) leading to L-optimality, where $\boldsymbol{W}$ is a symmetric nonnegative definite weighting matrix, $\lambda_{max}(\boldsymbol{M}^{-1})$ leading to E-optimality, where λ is an eigenvalue of $\boldsymbol{M}^{-1}$.

These are some interesting properties of a D-optimal design:

- *The optimum is invariant to transformations with non-singular Jacobians.* With these transformations, the location of the minimum does not change (its value however does change). An example is a scaling transformation. An transformation for which the minimum is not invariant is e.g. the transformation $\boldsymbol{t}$ from state space to the lower-dimensional goal space. E.g. the examples of Section 6.3.1 and 6.3.2 have a different goal, which gives rise to a different sensing strategy.

- *All units in the utility function are consistent.* This is merely a rephrasing of the previous property: if the units are consistent, a change of scale multiplies every term by the same scalar, and hence the minimum is invariant to this transformation. This is not the case with L- and E-optimal designs.

- *This utility function does establish a preference for 'poorly observed' directions,* despite the fact that reducing a singular value corresponding to a well observed direction has exactly the same effect on the volume of the ellipsoid as the same operation in poorly observed direction. Although the latter is true, this is not what happens, since the desired behaviour is obtained in combination with a Kalman filter. An observation applied to a poorly observed direction has a much larger effect on the volume of the ellipsoid as the same observation in a direction which has been observed many times before. Hence the former action is preferred by the utility function, which is the behaviour we are interested in.

- *The solution may correspond to a local minimum.* This is not too big a problem, since this concerns only one observation in the sequence of observations. Further investigation is needed to verify if this occurs frequently. Anyhow, the solution corresponding to this local minimum will probably still be better than a random sensing strategy.

- *The state covariance matrix being singular is no problem.* The utility function evaluates the determinant of the covariance matrix of the *goal location*, and not of the state estimate. Even though in our examples $\widehat{\boldsymbol{P}}$ is singular, $\widehat{\boldsymbol{P}}^g$ is not. In other words, $\boldsymbol{x}$ does not need to be a minimal state description of the object, while $\boldsymbol{x}^g$ does need to be a minimal description of the goal.

- *It is not clear whether the measurements thus planned can be considered as statistically independent.* This statistical independence among observations is a basic assumption of the KF. Here, we assume that this is true. However, one could object that each measurement location is determined using information from all previous ones (summarised in the KF estimate), and hence cannot be independent ...

- *It is not clear what the effect is of linearisation errors on the "optimality" of the actions.* It should however be clear that the optimality of the KF does not depend on the optimality of the planned actions. (Whaite and Ferrie 1997) limits the search for new viewpoints to a small region around the current viewpoint to limit linearisation

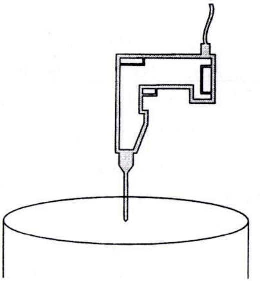

Figure 4: *Example:* drilling a hole in the centre of the top plane of a cylinder.

errors, which only is true if the surfaces to explore are sufficiently smooth (which is the case as objects are modelled using super-quadratics). Only in this case, the new chunks of information presented to the KF remain small, and the new estimate will remain close to the old estimate. However, there is no reason to have the same limitation here, as object surfaces can have edges where measurements vary dramatically; hence, even small changes in the position of the sensor can give rise to large variations in the measurement.

6 Results

Suppose the task is to drill a hole in the top plane of a bounded cylinder, Fig. 4. Measurements are taken with the ultrasonic sensor, Fig. 2. Section 6.1 first details the model of this cylinder. Section 6.2 shows how the measurement equation for the ultrasonic sensor is derived. Finally, Section 6.3 shows a planned sensing strategy for this example.

6.1 Object Model

The cylinder is modelled as a collection of two unbounded planes and one unbounded cylinder.

The parametrisation of a *plane*, recommended for optimal numerical stability (Anthony, Anthony, Cox, and Forbes 1991) is given by:

$$\left(\boldsymbol{U}_c \boldsymbol{a}(\alpha, \beta)\right)^T (\boldsymbol{x} - \boldsymbol{x}_c) - d = 0\,. \tag{14}$$

$(\boldsymbol{U}_c, \boldsymbol{x}_c)$ defines the locating frame of the plane in some reference frame, where $\boldsymbol{x}_c$ is the origin and $\boldsymbol{U}_c$ the rotation matrix of the frame relative to the reference frame. $|d|$ is the distance of the plane to $\boldsymbol{x}_c$. $\boldsymbol{a}(\alpha, \beta)$ is the orientation vector of the normal on the plane relative to the locating frame (which should be close to the z-axis of the locating frame for optimal stability), given by

$$\boldsymbol{a} = \begin{bmatrix} \cos(\beta) & 0 & \sin(\beta) \\ 0 & 1 & 0 \\ -\sin(\beta) & 0 & \cos(\beta) \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\alpha) & -\sin(\alpha) \\ 0 & \sin(\alpha) & \cos(\alpha) \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \tag{15}$$

$$= \begin{bmatrix} \cos(\alpha)\sin(\beta) \\ -\sin(\alpha) \\ \cos(\alpha)\cos(\beta) \end{bmatrix}, \tag{16}$$

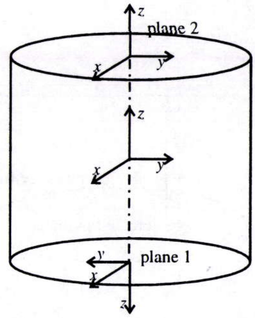

Figure 5: Relative position of the locating frames of the three features (the depicted position of the cylinder relative to these frames corresponds to $\boldsymbol{x}$ equal to a 10×1 zero vector).

or vector $\boldsymbol{a}$ is obtained by rotating a unit vector along the z-axis, first by an angle α around the x-axis, and then by an angle β around the old y-axis.

Thus, the state of the plane in the locating frame $(\boldsymbol{U}_c, \boldsymbol{x}_c)$ is described by $\boldsymbol{x} = \begin{bmatrix} \alpha & \beta & d \end{bmatrix}^T$.

The recommended parametrisation of a *cylinder* is given by :

$$||(\boldsymbol{x} - \boldsymbol{x}_c - \boldsymbol{U}_c\boldsymbol{x}_0) \times \boldsymbol{U}_c\boldsymbol{a}(\alpha, \beta)|| - r_0 = 0 \; . \tag{17}$$

Again, $(\boldsymbol{U}_c, \boldsymbol{x}_c)$ defines the locating frame of the cylinder in some reference frame, and $\boldsymbol{a}(\alpha, \beta)$ is the orientation vector of the centre line of the cylinder relative to the locating frame (which should be close to the z-axis of the locating frame for optimal stability), defined by Eqs. (16), and r_0 is the radius of the cylinder. $\boldsymbol{x}_0$ is a point satisfying $\boldsymbol{a}^T\boldsymbol{x}_0 = 0$. If $\boldsymbol{a}$ is close to $(0, 0, 1)$, this can be made explicit as

$$\boldsymbol{x}_0 = \begin{bmatrix} \xi \\ \eta \\ (-a_1\xi - a_2\eta)/a_3 \end{bmatrix} . \tag{18}$$

Thus, the state of the cylinder in the locating frame $(\boldsymbol{U}_c, \boldsymbol{x}_c)$ is described by $\boldsymbol{x} = \begin{bmatrix} \alpha & \beta & \xi & \eta \end{bmatrix}^T$.

Hence, the complete state description of the bounded cylinder is the 10×1 vector

$$\boldsymbol{x} = \begin{bmatrix} \alpha_c & \beta_c & \xi_c & \eta_c & \alpha_1 & \beta_1 & d_1 & \alpha_2 & \beta_2 & d_2 \end{bmatrix}^T . \tag{19}$$

Since a cylinder has only 5 dof (rotation around the axis is ambiguous) and $\boldsymbol{x}$ is 10-dimensional, 5 constraints exist between the 3 primitives: 4 linear constraints specifying that the plane normals and the centre line are parallel, and 1 nonlinear constraint specifying the distance between both planes. If the relative position of the locating frames is as shown in Fig. 5, then the 4 linear constraints are

$$\alpha_c - \alpha_1 = 0 \; , \tag{20}$$
$$\beta_c + \beta_1 = 0 \; , \tag{21}$$
$$\alpha_1 - \alpha_2 = 0 \; , \tag{22}$$
$$\beta_1 + \beta_2 = 0 \; . \tag{23}$$

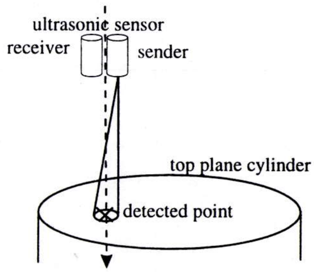

Figure 6: *Measurement with an ultrasonic sensor on a planar surface.*

The SCKF applies these linear constraints only once as perfect observations.

The distance l between the two planes is calculated as the distance between plane 2 and a point on plane 1. It is important to note that this distance has a *sign*, and is positive in the direction of the normal on plane 2. A point $\boldsymbol{p}$ on plane 1 is given by:

$$\boldsymbol{p} = \boldsymbol{x}_{c1} + \boldsymbol{a}_1 d_1 \ . \tag{24}$$

Substituting $\boldsymbol{p}$ in the equation of plane 2 yields the constraint relation $c(\alpha_1, \beta_1, d_1, \alpha_2, \beta_2, d_2) = 0$:

$$\boldsymbol{a}_2^T (\boldsymbol{x}_{c1} - \boldsymbol{x}_{c2} + \boldsymbol{a}_1 d_1) - d_2 - l = 0 \ . \tag{25}$$

The Jacobian $\nabla \boldsymbol{c}$ of this constraint is given by:

$$\nabla \boldsymbol{c} = \left[\begin{array}{cccccc} \frac{\partial c}{\partial \alpha_1} & \frac{\partial c}{\partial \beta_1} & \frac{\partial c}{\partial d_1} & \frac{\partial c}{\partial \alpha_2} & \frac{\partial c}{\partial \beta_2} & \frac{\partial c}{\partial d_2} \end{array} \right] . \tag{26}$$

Calculation of these partial derivatives is a straightforward but tedious job. Again, the SCKF applies this constraint as a normal observation (replace $\boldsymbol{H}_i$ by $\nabla \boldsymbol{c}_i$ in the KF eqs.), but now smoothly to avoid a bias on the estimate due to linearisation errors: the SCKF applies this constraint as a set of linearised constraints, of which the variance is artificially increased; The application of these linearised constraints is interlaced with "normal" measurements, until a stop criterion is satisfied (De Geeter et al. 1997).

6.2 Measurement Equations

The distance measured by the ultrasonic (US) sensor is calculated as the intersection of the line-of-sight of the sensor and the surface, e.g. plane 2, Fig. 6.2.

The line-of-sight is determined by the position of the robot, and does not need to be estimated. Therefore, the numerical stability of this equation is not important. The following parametrisation is chosen for this line:

$$\boldsymbol{x} = \boldsymbol{x}_{US} + \boldsymbol{a}_{US}\, t \ , \tag{27}$$

where $\boldsymbol{x}_{US}$ is the position of the US sensor, $\boldsymbol{a}_{US}$ is a unit vector in the direction of the line-of-sight of the US sensor, and hence, t is the distance on the line between $\boldsymbol{x}$ and $\boldsymbol{x}_{US}$.

Thus, the scalar measurement $z = h(\alpha_2, \beta_2, d_2)$ is equal to the value of t at the intersection point of the line and the plane. Substituting the expression for the line-of-sight, Eq. (27), in the equation of the plane, Eq. (14), and expliciting for t yields:

$$t = \left(d_2 + \boldsymbol{a}_2^T (\boldsymbol{x}_{c2} - \boldsymbol{x}_{US})\right) / \left(\boldsymbol{a}_2^T \boldsymbol{a}_{US}\right) \ . \tag{28}$$

The Jacobian $\nabla\boldsymbol{H}$ is calculated as follows:

$$\nabla\boldsymbol{H} = \left[\begin{array}{ccc} \frac{\partial h}{\partial \alpha} & \frac{\partial h}{\partial \beta} & \frac{\partial h}{\partial d} \end{array} \right] , \qquad (29)$$

where again the calculation of the partial derivatives is tedious but straightforward.

The equation for an US measurement plane 1 and on the cylinder surface can be found in a similar way.

6.3 A Planned Sensing Strategy

This section compares two examples: in the first, the goal is to drill a hole in the top plane, and in the second, the goal is simply to find the centre point of the top plane. The difference is that in the second example, the orientation of the top plane is not important.

Both sensing strategies are generated under the same conditions. The initial estimate is :

$$\hat{\boldsymbol{x}}_0 = \left[\begin{array}{cccccccccc} 0.1 & 0.1 & 0 & 0 & 0 & 0 & 0.1 & 0 & 0 & 0 \end{array} \right]^T \qquad (30)$$

$$\widehat{\boldsymbol{P}}_0 = \mathrm{diag}(\; 0.03 \;\; 0.03 \;\; 0.01 \;\; 0.01 \;\; 0.01 \;\; 0.01 \;\; 0.03 \;\; 0.01 \;\; 0.01 \;\; 0.01 \;) \qquad (31)$$

The standard deviation of the noise on the US distance measurement is equal to .01m. The cylinder has a radius of 1m and a length of 2m.

In these simulations, the MHT consists of only one node on each level (i.e. the measurement-feature correspondence is known), which makes the figures more easy to interpret.

6.3.1 Example 1: Drilling a Hole in the Top Plane

In this example, both the position of the centre point of the top plane and the orientation of the normal of the top plane are important. Hence the goal is described by $\boldsymbol{x}^g = \left[\begin{array}{ccccc} x^g & y^g & z^g & \alpha^g & \beta^g \end{array} \right]^T$. The first three rows of transformation $\boldsymbol{t}$ are found by first calculating the intersection between the top plane and the centre line of the cylinder, and second, calculating the Jacobian of this transformation. The last two rows simply state that $\alpha^g = \alpha_2$ and $\beta^g = \beta_2$.

Figure 7 shows the expected value of the change in the utility caused by an action, on the different primitives of the cylinder, for the first 5 time steps. This utility gain is calculated as

$$\Delta ua_j(\boldsymbol{v}_j)_{i+1} = ua_j(\boldsymbol{v}_j)_{i+1} - u_i \; , \qquad (32)$$

where u_i is the current value of the utility function, $ua_j(\boldsymbol{v}_j)_{i+1}$ is the predicted value of the utility on the next time step, given action $a_j(\boldsymbol{v}_j)$, calculated as in equation 13. The optimum in this case is at the minimum of this difference (or at the largest absolute value).

The resulting measurements are as follows, Fig. 7:

Measurement 1: The SCKF did not yet apply any constraints, and hence the 3 primitives behave completely independent. Therefore, the utility function is zero over the whole bottom plane. The best measurement is somewhere on the top edge of the cylinder surface (the choice is arbitrary since the utility is equal for all points on this top edge).

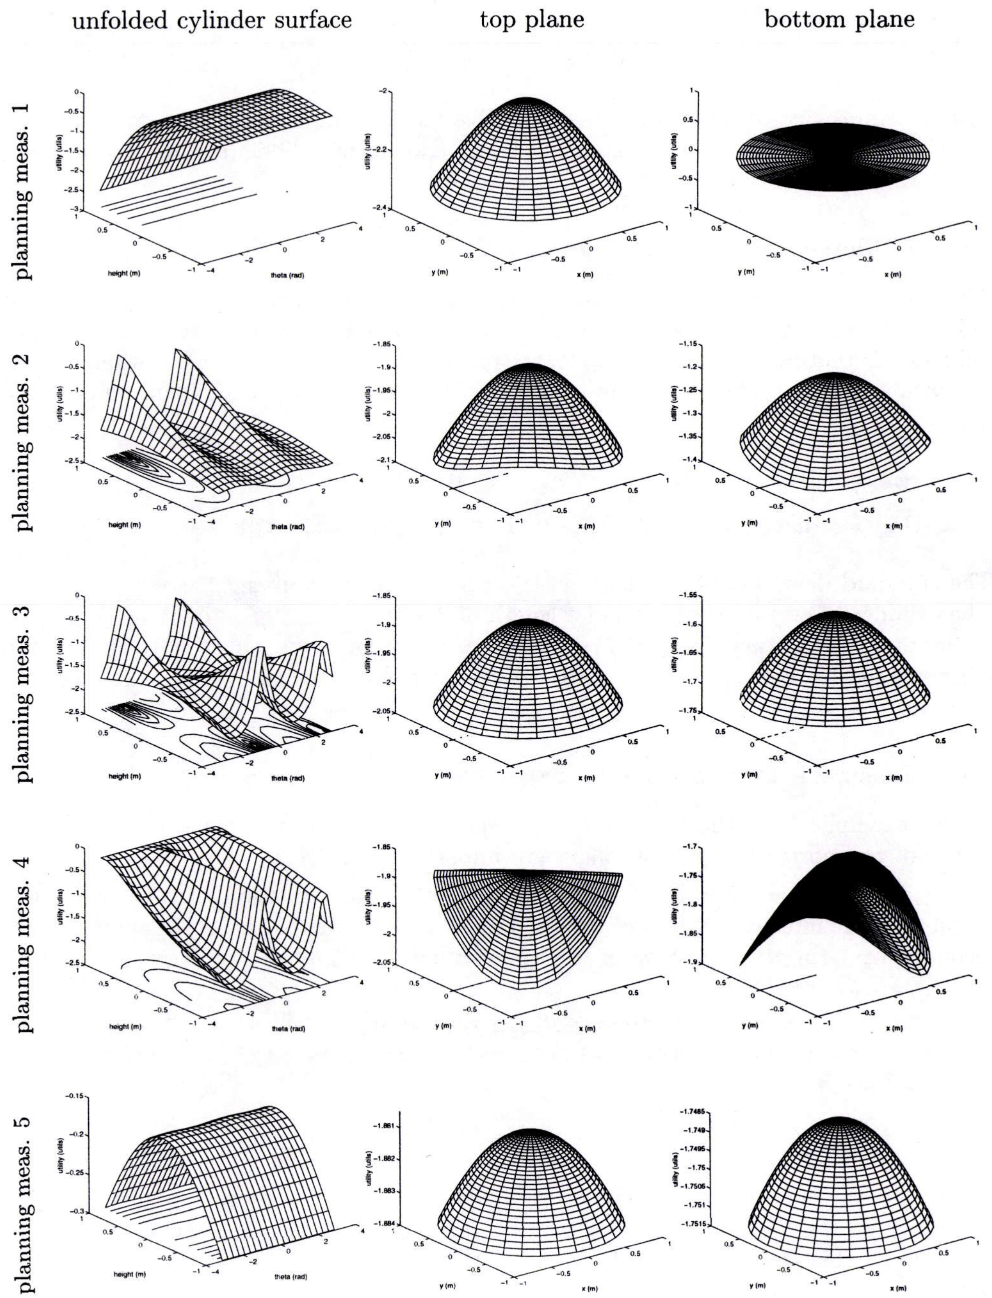

Figure 7: Drilling a hole in the centre of the top plane. The expected value of the change in utility, eq. (32), for a measurement on (from left to right), the unfolded cylinder surface, the top plane and the bottom plane, for measurements (from top to bottom) 1 to 5.

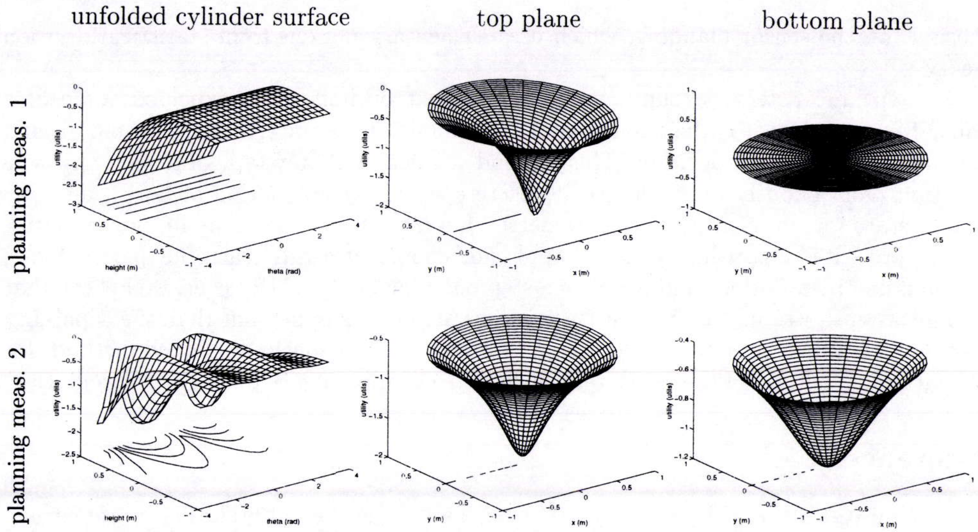

Figure 8: Finding the position of the centre point of the top plane. The expected value of the change in utility, eq. (32), for a measurement on (from left to right), the unfolded cylinder surface, the top plane and the bottom plane, for measurements (from top to bottom) 1 and 2.

Measurement 2: The effect of this measurement is clearly visible: it is not interesting to go back to the same point again. The best measurement is again on the top edge of the plane, but now 90 degrees shifted relative to the first measurement. Note that now, due to the application of the constraints, the utility on the bottom plane is no longer zero: the cylinder now behaves as a rigid body.

Measurement 3: on the bottom edge of the cylinder.

Measurement 4: on the bottom edge of the cylinder.

Measurement 5: on the edge of the top plane. Note that the utility gain on the cylinder surface is now about an order of magnitude smaller than the utility gain on the top plane.

6.3.2 Example 2: Finding the Position of the Centre Point of the Top Plane

Now, the goal is $\boldsymbol{x}^g = \begin{bmatrix} x^g & y^g & z^g \end{bmatrix}^T$. The Jacobian of the transformation $\boldsymbol{t}$ is found in a similar way as above. Fig. 8 shows the first two measurements thus planned. The most important difference is the dip in the utility function of the top plane, around the centre point. This is because the orientation of the plane is not important. Still, the first measurement is taken on the top edge of the cylinder surface, since this is the only way to determine x^g and y^g.

7 Conclusion

This paper presents an integrated Bayesian solution to the problem of object location estimation, object recognition and planning sensor actions under uncertainty. The emphasis is on the sensor planning, which uses elementary notions from statistical decision theory.

The optimal action is found as the action that optimises the expected value of a utility function, which is the logarithm of the volume of the uncertainty ellipsoid around the estimate of a target position. This method can deal with any type of target, any type of action (continuously parametrised, a discrete set or a combination of both), and any type of state (again, continuously parametrised, a discrete set or a combination of both).

The presented algorithm is easy to apply and computationally tractable, partly due to the Gaussian assumptions and to the one-step planning horizon. However, we believe that the limitations introduced by these assumptions are far less important than the capability of on-line planning, simplicity, and integration with the previously developed methods for estimation, resolution of the measurement-feature correspondence and object recognition.

References

Anthony, G. T., H. M. Anthony, M. Cox, and A. Forbes (1991). The parametrisation of geometric form. Technical Report EUR 13517, Commission of the European Communities.

Bajcsy, R. (1988, August). Active perception. *Proceedings of the IEEE 76*(8).

Bar-Shalom, Y. and X.-R. Li (1993). *Estimation and Tracking: Principles, Techniques and Software.* Artech House, Norwood.

Barraquand, J. and P. Ferbach (1995). Motion planning with uncertainty, the information space approach. In *Proceedings IEEE International Conference on Robotics and Automation*, pp. 1341–1348.

Borghi, G. and V. Caglioti (1995, September). Optimum explorations in the self-localisation of mobile robots in polygonal environments. In *Proceedings International Conference on Advanced Robotics*, Sant Filiu de Guixols, Spain, pp. 993–997.

Bouilly, B., T. Siméon, and R. Alami (1995). A numerical technique for planning motion strategies of a mobile robot in presence of uncertainty. In *Proceedings IEEE International Conference on Robotics and Automation*, pp. 1327–1332.

Cook, D. J., P. Gmytrasiewicz, and L. B. Holder (1996, October). Decision-theoretic cooperative sensor planning. *IEEE Transactions on Pattern Analysis and Machine Intelligence 18*(10), 1013–1023.

Das, S., G. Beni, and S. Hackwood (1992, July). An optimal sensing strategy for mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Raleigh, NC, pp. 393–398.

De Geeter, J., J. De Schutter, H. Van Brussel, and M. Decréton (1997, April). A smoothly constrained Kalman filter. Technical report, K.U.Leuven, Dept. of Mechanical Engineering.

De Geeter, J., H. Van Brussel, J. De Schutter, and M. Decréton (1996). Recognising and locating objects with local sensors. In *Proceedings IEEE International Conference on Robotics and Automation*, Minneapolis, pp. 3478–3483.

De Geeter, J., H. Van Brussel, J. De Schutter, and M. Decréton (1997, September). Local world modelling for teleoperation in a nuclear environment using a Bayesian multiple hypothesis tree. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Grenoble, France.

Decréton, M. (1995, March). Distance sensing in nuclear remote operation. *Measurement 15*(1), 43–51.

Duncan Luce, R. and H. Raiffa (1957). *Games and decisions: introduction and critical survey.* New York: John Wiley & Sons.

Erdmann, M. (1995, October). Understanding action and sensing by designing action-based sensors. *International Journal of Robotics Research 14*(5), 483–509.

Gelb, A. (1974). *Applied Optimal Estimation.* The Analytic Sciences Corporation.

Goodwin, G. C. and R. L. Payne (1977). *Dynamic System Identification.* Academic Press.

Hager, G. D., S. P. Engelson, and S. Atiya (1993). On comparing statistical and set-based methods in sensor data fusion. In *Proceedings IEEE International Conference on Robotics and Automation*, pp. 352–358.

Kalman, R. E. (1960, March). A new approach to linear filtering and prediction problems. *Transactions of the ASME, Journal of Basic Engineering*, 35–45.

Kristensen, S. (1995, July). Sensor planning with Baysesian decision theory. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Pisa, Italy, pp. 237–248.

LaValle, S. M. (1995). *A game-theoretic framework for robot motion planning.* Ph. D. thesis, Univ. of Illinois at Urbana-Champaign.

Mehra, R. K. (1974). Optimal inputs for linear system identification. *IEEE Transactions on Automatic Control 19*(3), 192–200.

Page, L. A. and A. C. Sanderson (1995). Robot motion planning for sensor-based control with uncertainties. In *Proceedings IEEE International Conference on Robotics and Automation*, pp. 1333–1340.

Sorenson, H. W. (1985). *Kalman Filtering: Theory and Application.* IEEE Press.

Subrahmonia, J., D. B. Cooper, and D. Keren (1996, May). Practical reliable Bayesian recognition of 2D and 3D objects using implicit polynomials and algebraic invariants. *IEEE Transactions on Pattern Analysis and Machine Intelligence 18*(5), 505–519.

Whaite, P. and F. P. Ferrie (1997, March). Autonomous exploration: driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence 19*(3), 193–205.

Zarrop, M. B. (1979). *Optimal experiment design for dynamic system identification*, Volume 21 of *Lecture notes in control and information science.* Berlin: Springer-Verlag.