Experiments in Control of Sensory Perception for Robotic Systems

Tomasz Celinski and Brenan McCarragher
Department of Engineering
Faculty of Engineering and Information Technology
The Australian National University
Canberra, Australia
Fax: int + 61 6 249 0506
E-mail: {tomasz,brenan}@faceng.anu.edu.au

Abstract

The paper presents a new approach to management of multiple sensors and perception algorithms in a multi-sensor robotic system. The approach involves real-time selection of process monitors by a Sensory Perception Controller. The selection is based on the minimization of the expected cost of perception with constraints on the uncertainty of perception. The effectiveness and usefulness of the approach is evaluated through experiments involving a range of sensing modalities which may typically be encountered in robotic applications

1 Introduction

Sensory perception is one of the essential elements of an intelligent autonomous system [Brooks, 1986]. Autonomous systems operating in uncertain, dynamically changing environments have large, time varying information requirements. The requirements necessitate the use of a wide range of sensors for data acquisition, as well as data processing algorithms which reduce raw sensory data to a form suitable for various levels of modeling and decision making.

The extensive perceptual requirements of robotic systems coupled with the need for efficient operation can only be satisfied through intelligent perception. We propose that an intelligent perception system should have the following characteristics:

- Multiple Sensors and Perception Algorithms: The system will have to be equipped with multiple sensors and perception algorithms. The array of sensors and algorithms will have to be sufficiently wide to cover the required range of operating conditions.
- Adaptability and Fault-Tolerance: The operation of the perception system should be adaptable to environmental conditions and tolerant to faults within the perception system. When some of the sensors and/or algorithms perform poorly or fail the system should either adapt the relevant algorithms to improve performance, or other sensors and/or algorithms should be employed. The overall performance should either not be affected at all, or it should degrade gracefully.
- Extensibility: The addition of new sensors and algorithms to the system should not require any significant changes to the overall system architecture. The overall system performance should not decrease as a result of an addition.
- Efficiency: Perceptual resources should be used in a way which ensures efficient delivery of information which is relevant and of sufficient quality with respect to the overall state of the system.

The purpose of this paper is to present an approach to sensory perception which satisfies the above requirements. The approach involves a Sensory Perception Controller and the work falls into the areas of sensor management, planning and control.

Early work on control of sensory perception, such as [Meier et al., 1967; Athans, 1972] considered off-line planning of sensing activities. Such approaches are not able to adapt the plans to changes in the task and in the environment. They are therefore not suitable for autonomous systems operating in unpredictable, dynamic environments. Our approach overcomes this limitation.

An approach to planning sensing strategies on-line is presented in [Hutchinson and Kak, 1989]. The system automatically proposes a sensing operation and then determines the maximum ambiguity, measured by entropy, which might remain in the world description if the proposed sensing operation is applied. The sensing operation which is actually carried out is the one which minimizes the ambiguity. The approach concentrates on the problems of object recognition and pose estimation.

A recent approach to the control of sensory perception can be found in [Hovland and McCarragher, 1997]. The approach considers the control of sensory perception in a real-time discrete event control system. The aim of the

sensory control is to efficiently collect perceptual information for the purpose of high confidence discrete event recognition. The sensory perception controller is formulated as a sequential Markov decision problem and the decision problem is solved using stochastic dynamic programming.

The work presented in this paper builds on [Celinski and McCarragher, 1998a; 1998b] and is significantly different from earlier approaches. It centers around a sensory perception controller (SPC) which performs real-time selection of process monitors (a process monitor is a combination of sensors and a perception algorithm). The SPC receives a request for sensory data formulated in terms of information primitives, which correspond to the types of information used by the system. The SPC then proceeds to select and activate process monitors, aiming to satisfy the request in the most efficient manner.

2 Exchange of Information

The specification of information requirements to the SPC is based on the concept of an information primitive:

Definition 2.1 (Information Primitive) *An information primitive is any dynamic state variable, discrete event state variable, model parameter, or task parameter which has a degree of uncertainty during task execution.*

An information primitive (IP) will be denoted by η and the collection of all IPs by $\mathcal{I}$. The values which η may assume will be restricted to be numeric.

Information primitives form one of the components of an information request. A request must also specify any constraints on the information which the SPC is to deliver. A useful and important class of constraints are uncertainty constraints, which allow specification of upper bounds on the uncertainty of the estimates produced by the SPC. We define an information request as:

Definition 2.2 (Information Request) *An information request from an information consumer, h, to the SPC is given by $\mathcal{R}^h = (\mathcal{I}^h, \Delta^h)$ where $\mathcal{I}^h$ is a set of required information primitives, and $\Delta^h = \{\delta^h_\eta : \eta \in \mathcal{I}^h\}$ is a set of uncertainty constraints.*

3 Modeling of Perceptual Capabilities

The system and the environment are observed by sensors and the sensory data is processed using appropriate algorithms. The combination of sensing and data processing will be referred to as process monitoring:

Definition 3.1 (Process Monitor) *A process monitor is a combination of one or more sensors and sensory data processing algorithms. Each monitor m outputs either a set of probability densities $\hat{\mathcal{I}}_m = \{f_\eta(\eta)) : \eta \in \mathcal{I}_m\}$, or a set of confidence values $\hat{\mathcal{I}}_m = \{q_\eta : \eta \in \mathcal{I}_m\}$, expressing the knowledge which m acquired about the information primitives $\eta \in \mathcal{I}_m$.*

The confidence values referred to in the above definition are real numbers in the range $[0, 1]$. A value of 0 indicates a complete lack of confidence (no knowledge), while a 1 indicates complete confidence (full knowledge).

The performance of a process monitor may be described in terms of its uncertainty and cost characteristics. The uncertainty measures for a monitor m may be either calculated based on the density $f_\eta(\eta)$ output by m, or based on confidence values output by m. The cost characteristic reflects the amount of resources consumed by a single observation. The uncertainty and cost characteristics constitute a process monitor model:

Definition 3.2 (Process Monitor Model)

(a) For monitor m outputting probability densities: The process monitor model for m is the triple $(\mathcal{I}_m, \mathcal{F}_m, c_m(\underline{\alpha}_m))$. $\mathcal{I}_m$ is the set of primitives which m is able to monitor; $\mathcal{F}_m = \{f_{m,\eta}(\eta; \underline{\alpha}_m) : \eta \in \mathcal{I}_m\}$ is a set of probability densities expressing the uncertainty characteristics, where $\underline{\alpha}_m$ is a parameter vector reflecting the operating conditions and $c_m(\underline{\alpha}_m)$ is the cost of using m.

(b) For monitor m outputting confidence values: The process monitor model for m is the triple $(\mathcal{I}_m, \mathcal{Q}_m, c_m(\underline{\alpha}_m))$. $\mathcal{I}_m$ is the set of primitives which m is able to monitor; $\mathcal{Q}_m = \{q_\eta(\underline{\alpha}_m) : \eta \in \mathcal{I}_m\}$ is a set confidence functions expressing the uncertainty characteristics, where $\underline{\alpha}_m$ is a parameter vector reflecting the operating conditions and $c_m(\underline{\alpha}_m)$ is the cost of using m.

The parameters $\underline{\alpha}_m = [\alpha_1, ..., \alpha_r]$ are a subset of the information primitives. They define the monitor's dependence on the system and the environment. For example, vision depends on lighting conditions (l) and the camera-target distance (d), giving $\underline{\alpha}_m = [l, d]$. Vision is only useful over a limited range of illuminations and distances.

4 Perception Control Algorithm

The perception control algorithm determines which of the process monitors should be used at each stage of task execution. The algorithm is illustrated in Figure 1

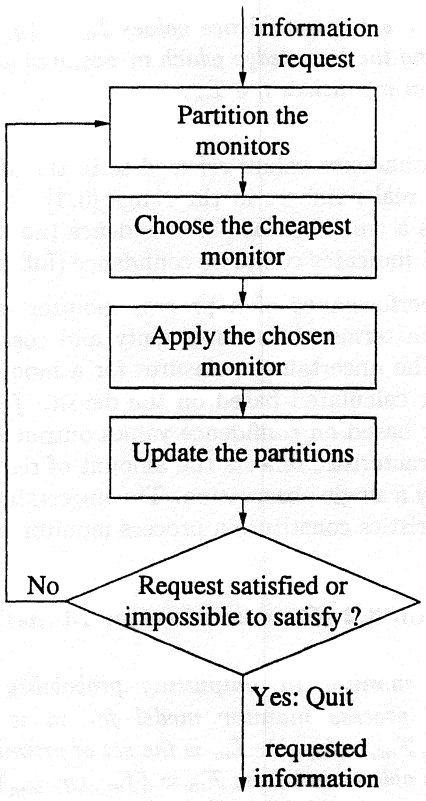

Figure 1: The sensory perception control algorithm

The first step of the algorithm is to partition the available process monitors into two sets: $\mathcal{M}_{<\Delta}$ and $\mathcal{M}_{\geq\Delta}$. The first set, $\mathcal{M}_{<\Delta}$, contains monitors whose observation uncertainty is expected to fall within the specified bounds Δ (the bound is part of the information request received by the SPC), the second set contains all the remaining process monitors.

The cost of monitoring is used as a criterion for deciding which monitor from the partition $\mathcal{M}_{<\Delta}$ should be activated. The chosen monitor is:

$$m^* = \underset{m \in \mathcal{M}_{<\Delta}}{\operatorname{argmin}}\{c_m\} \quad (1)$$

where c_m is an expected cost.

If $\mathcal{M}_{<\Delta}$ is empty then $\mathcal{M}_{\geq\Delta}$ is used instead. The chosen monitor is then applied and it delivers new information about a primitive η. If this information has uncertainty smaller than the bound δ_η then all monitors of η are removed from $\mathcal{M}_{<\Delta}$ and $\mathcal{M}_{\geq\Delta}$, otherwise only m^* is removed. The algorithm terminates if all the requested primitives have been adequately observed, or there are no process monitors left in the sets $\mathcal{M}_{<\Delta}$ and $\mathcal{M}_{\geq\Delta}$.

The above discussion provides an overview of the Perception Control Algorithm. A detailed discussion may be found in [Celinski and McCarragher, 1998a; 1998b].

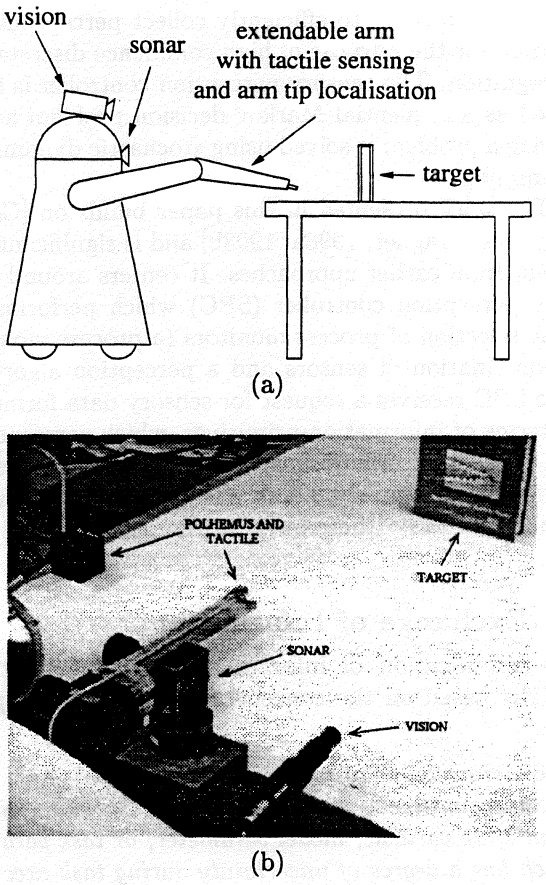

Figure 2: (a) The experimental scenario (b) Actual setup

5 Experimental Evaluation of the SPC

We now turn to an experimental demonstration of the effectiveness of control of sensory perception. The demonstrations will be carried out on a task involving object recognition and localization. Recognition and localization tasks were chosen because they form a significant class of perception tasks for robots.

The experimental scenario is shown in Figure 2 (a) and the actual setup is shown in Figure 2 (b). The scenario involves determining of (1) the target object type and (2) the distance to the object. The objects are members belonging to a known set. The set of objects used in the experiments consisted of three different books.

The perceptual system consists of a number of sensors and data processing algorithms. The sensors included vision, a combination of tactile sensing and electromagnetic position sensing on an extendable platform, as well as sonar. The sensors were combined with data processing algorithms giving the following seven process monitors:

- $m_{recog1}, m_{recog2}, m_{recog3}$: target classifiers based on

Monitor	Cost
m_{recog1}	0.94sec
m_{recog2}	0.94sec
m_{recog3}	0.94sec
$m_{tactile}$	0.1875 sec/cm
m_{sonar}	1.03sec
m_{edge}	4.17sec
m_{light}	0.23sec

Table 1: Monitoring costs. The costs of m_{recog1}, m_{recog2}, m_{recog3}, m_{sonar}, m_{edge}, and m_{light} were fixed. The cost of $m_{tactile}$ increased with target distance at a rate of $0.1875 sec/cm$.

binary vision and template matching. The monitors rely on natural scene illumination and their performance degrades with decreasing lighting as well as increasing target distance. They differ in the light level at which their performance is optimal. The experiments were conducted in a room in which lighting could be set at one of five possible levels

- $m_{tactile}$: a distance monitor based on a Polhemus 3SPACE FASTRAK electromagnetic position measurement instrument and a contact sensor. The Polhemus instrument provides measurements of the relative position between a receiver and a transmitter (see [Polhemus Inc, 1993] for details about the sensor). The receiver is located on a moving platform which may be extended to touch a target object. Distance measurements are taken when the arm is in contact with a target. The uncertainty of the monitoring increases with increasing distance between the transmitter and the receiver. The cost of using the monitor increases with measured distance.

- m_{sonar}: a sonar-based distance monitor. The monitor has a poor resolution but is very cheap to use.

- m_{edge}: a vision based distance monitor. The monitor obtains distance estimates based on edges detected within the image. Knowledge of the location of the surface on which target objects are located is combined with edges corresponding to the junctions between target objects and the surface to give distance estimates.

- m_{light}: a vision-based light intensity monitor.

The seven process monitors were applied under a range of operaing conditions and their monitoring uncertainty was captured in uncertainty and cost models (see Definition 3.2). An example of such an uncertainty model is given in Figure 3. The process monitor cost models are summarized in Table 1.

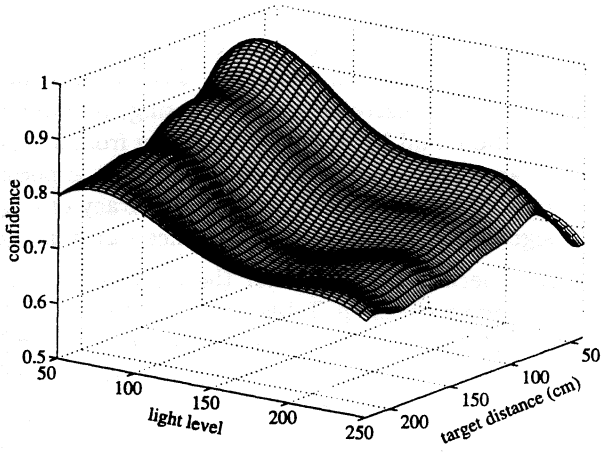

Figure 3: The uncertainty model of m_{recog1}

5.1 Experiments

The SPC was presented with the following information requests:

$$\begin{aligned} \mathcal{R}^1 &= (\{t,\ d\},\ \{\delta_t = 0.2,\ \delta_d = 5cm\}) \\ \mathcal{R}^2 &= (\{t,\ d\},\ \{\delta_t = 0.25,\ \delta_d = 6.5cm\}) \\ \mathcal{R}^3 &= (\{t,\ d\},\ \{\delta_t = 0.15,\ \delta_d = 3.5cm\}) \end{aligned}$$

The information requests are formulated according to Definition 2.2. Request $\mathcal{R}^1$ indicates that the SPC should determine the target type (t) with recognition uncertainty of no more than 0.2 (recognition confidence of at least $1 - \delta_t = 0.8$) and target distance (d) with standard deviation of at most $5cm$ under a normal distribution assumption. Requests R^2 and R^3 may be interpreted in a similar way.

The information gathering was performed under the following experimental conditions:

- E1: In this experiment there was very little prior knowledge about the distance to and type of target, and the light level The experiment was not affected by any disturbances.

- E2: Prior knowledge of the type and distance to the target was very limited (uniform density), but the light level was given by a known Gaussian density with mean 249 and standard deviation of 2.5. The experiment was affected by a malfunction of the sonar.

- E3: Prior knowledge of target distance was again limited, but light level was believed to be known at 160. The actual light level was affected by temporary shadows and was approximately 54. The experiment was also affected by a sonar malfunction.

- E4: This experiment assumed very little prior knowledge about the distance to the target or the light level (uniform densities). The experiment was affected by extremely poor lighting, a sonar malfunction, and by EM noise radiating from a vacuum cleaner operating in the vicinity of the experiment (the EM noise made the actual accuracy of $m_{tactile}$ significantly different from its uncertainty model).

The experimental results for the information request $\mathcal{R}^1$ are summarized in Table 2(a) and Figure 4. The experiments were also repeated for requests $\mathcal{R}^2$ and $\mathcal{R}^3$ Tables 2(b) and (c) and Figure 5 summarise the SPC's behaviours resulting from the different requests.

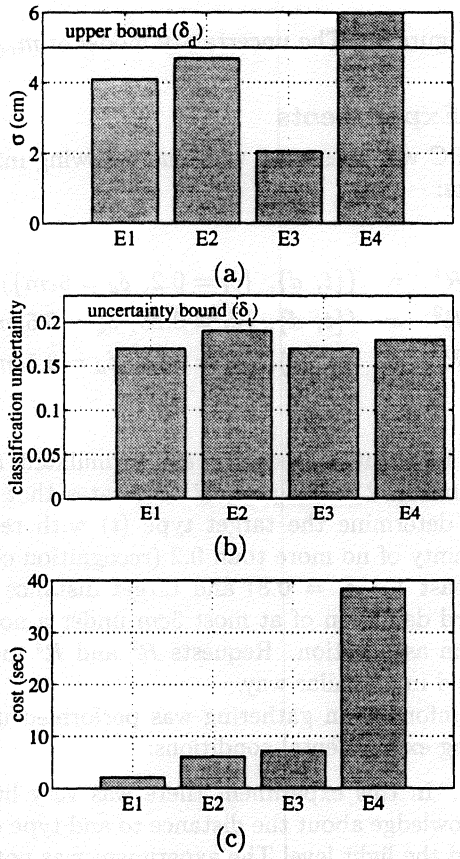

Figure 4: Experimental results for request $\mathcal{R}^1$ (a) standard deviations of distance estimates (b) target classification uncertainties (c) cost of perception

5.2 Discussion

The experiments highlighted the main characteristics of the Sensory Perception Controller.

The SPC dynamically *controlled* multiple sensors and perception algorithms (process monitors). The sequences of monitoring operations obtained in the experiments, shown in Table 2, were not fixed and were not determined *a priori*. The sequences were a result

Experiment	Monitor sequence
E1	m_{sonar}, m_{light}, m_{recog1}
E2	m_{recog3}, m_{sonar}
E3	m_{recog2}, m_{sonar}, m_{edge}
E4	m_{sonar}, m_{light}, m_{recog2}, m_{edge}

(a)

Experiment	Monitor sequence
E1	m_{sonar}, m_{light}, m_{recog1}
E2	m_{recog3}, m_{sonar}, m_{edge}
E3	m_{recog2}, m_{sonar}, m_{edge}, m_{light}, m_{recog_1}
E4	m_{sonar}, m_{light}, m_{recog2}, m_{edge}, $m_{tactile}$

(b)

Experiment	Monitor sequence
E1	m_{light}, m_{recog1}, m_{edge}, $m_{tactile}$
E2	m_{recog3}, m_{edge}
E3	m_{recog2}, m_{light}, m_{recog_1}, m_{edge}
E4	m_{light}, m_{recog2}, m_{edge}, $m_{tactile}$, m_{sonar}

(c)

Table 2: Sequences of monitoring operations performed by the SPC in response to (a) $\mathcal{R}^2$, (b) $\mathcal{R}^1$ and (c) $\mathcal{R}^3$. The order of requests ($\mathcal{R}^2$, $\mathcal{R}^1$, $\mathcal{R}^3$) reflects the tightness of the constraints ($\mathcal{R}^2$ least tight, $\mathcal{R}^3$ tightest). The varying sequences demonstrate the dynamic, adaptable, and fault-tolerant nature of the SPC.

of a closed-loop decision making process which occurred on-line and took into account information about the expected performance characteristics of the process monitors, the operating conditions, as well as the actual performance of the process monitors.

In cases where, due to disturbances or malfunctions (such as in E3), the selected monitors did not deliver information of sufficiently high quality, the SPC consulted other, more expensive monitors. The SPC was thus able to provide adaptable and fault-tolerant perception with gradual and graceful degradation of performance reflected by increased perception costs. This

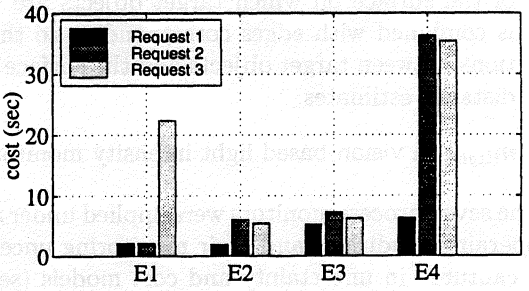

Figure 5: Comparison of the costs of monitoring for requests with different uncertainty constraints.

behavior satisfied the guidelines for intelligent perception outlined in Section 1. The SPC failed to satisfy information requests only if the uncertainty constraints of the requests exceeded the available perceptual capabilities (but it did return the best estimates which it was able to obtain). The gradual degradation in performance is shown in Figure 4(c). The cost of perception was lowest in experiment E1, where each of the selected process monitors performed well. In E2 monitor m_{sonar} was selected but failed to deliver a sufficiently good distance estimate because the target was far away from it. The request was satisfied through monitor m_{edge}, but at extra cost. Experiment E3 involved a malfunction of m_{sonar} as well as misleading prior belief about the lighting conditions. The SPC recovered from both these difficulties, but at relatively high cost. The last experiment, E4, presented an impossible challenge because the request exceeded the perceptual capabilities of the system. The SPC was able to recognize the target with sufficiently high confidence, but was unable to determine the distance even though it used all of the relevant process monitors. This outcome is shown in Figure 4(a), where σ_d for E4 exceeds the requested bound.

The perception achieved by the SPC was not only adaptable and fault-tolerant but also *efficient*. The algorithm selected the cheapest process monitors which satisfied the constraints of the requests. Efficiency was further improved by seeking and using low-cost information which was not directly relevant to the requests, but facilitated better decision making. For example, in experiment E1 the light level was determined (at a small cost) before deciding whether to use m_{recog1}, m_{recog2}, or m_{recog3}. This allowed the SPC to avoid the significant cost of using inappropriate process monitors.

The flexible, dynamic nature of the SPC allowed it to handle performance levels which were specified online. In the experiments the SPC was able to deal with three information requests ($\mathcal{R}^1$, $\mathcal{R}^2$ and $\mathcal{R}^3$) which imposed different constraints on the information gathering task. The perception control lead to three different sets of monitoring sequences, shown in Table 2, which reflected the constraints imposed by the requests. The different constraints also had a significant effect on the cost of monitoring, as illustrated in Figure 5. The figure shows a trend of increasing cost with increasing tightness of constraints, reflecting the higher perceptual challenge posed by tight constraints.

The SPC satisfies the requirements of an intelligent perception system outlined in Section 1. This, combined with the relative simplicity of the scheme, make it well suited for implementation in robotic systems.

6 Conclusion

The article introduces the Sensory Perception Controller as a new technique for dynamically managing the perceptual resources of robotic systems. The SPC has four main characteristics: (1) it allows the desired levels of performance to be specified on-line through information requests (2) it efficiently satisfies the information requests by minimizing the expected cost of perception (3) it provides a significant degree of robustness in cases where monitors fail to deliver sufficient information and (4) it uses information obtained on-line to improve its decision making. These four characteristics, combined with the relative simplicity of the scheme make it well suited for implementation in robotic systems.

References

[Athans, 1972] Athans, M. (1972). On the determination of optimal costly measurement strategies for linear stochastic systems. *Automatica*, 8:397–412.

[Brooks, 1986] Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23.

[Celinski and McCarragher, 1998a] Celinski, T. and McCarragher, B. (1998a). Determining the value of monitoring for dynamic monitor selection. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automat.*

[Celinski and McCarragher, 1998b] Celinski, T. and McCarragher, B. (1998b). Maximising the value of information primitives for dynamic monitor selection. In *Proceedings of the World Automation Congress WAC 98.*

[Hovland and McCarragher, 1997] Hovland, G. and McCarragher, B. (1997). Dynamic sensor selection for robotic systems. In *Proc. 1997 IEEE Int. Conf. Robotics and Automat.*, pages 272–277.

[Hutchinson and Kak, 1989] Hutchinson, S. and Kak, A. (1989). Planning sensing strategies in robot work cell with multi-sensor capabilities. *IEEE Trans. Robot. Automat.*, 5:765–783.

[Polhemus Inc, 1993] Polhemus Inc. (1993). *3SPACE FASTRAK Users Manual Revision F.*

[Meier et al., 1967] Meier, L., Peschon, J., and Dressler, R. (1967). Optimal control of measurement subsystems. *IEEE Transactions on Automatic Control*, 12:528–536.