

On the advantages of foveal mechanisms for active stereo systems in visual search tasks

Rui Figueiredo*, Alexandre Bernardino*, José Santos-Victor*, Helder Araújo†

*Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Portugal

†Institute for Systems and Robotics, Universidade de Coimbra, Portugal

{ruifigueiredo, alex, jasv}@isr.tecnico.ulisboa.pt, helder@isr.uc.pt

Abstract—In this work we study how information provided by foveated images sampled according to the log-polar transformation can be integrated over time in order to build accurate world representations and accomplish visual search tasks in an efficient manner. We focus on a specific visual information modality – depth – and on how to store it in a flexible memory structure. We propose a probabilistic observational model for a stereo system that relies on the Unscented Transform in order to propagate uncertainty in stereo matching, due to spatial quantization in the retina, to the 3D Cartesian domain. Probabilistic depth measurements are integrated in a novel Sensory Ego-Sphere whose topology can be biased with foveal-like distributions, according to the autonomous agent short-term tasks and goals. Furthermore, we investigate an Upper Confidence Bound (UCB) algorithm for the task of simultaneously finding the closest object to the observer (visual search) and learning the surrounding environment 3D map (mapping). The performance of task execution is assessed both with a foveated log-polar sensor and a classical uniform one. The advantage of foveal vision and custom ego-sphere representations are illustrated in a series of experiments with a realistic simulator.

Keywords—*Stereoscopic Vision, Foveal Vision, Active Vision, Sensory Ego-Sphere*

I. INTRODUCTION

The spatial organization of the photo-receptors in the human retina is not uniform. Visual acuity is highest at the fovea and declines monotonically to the periphery with increasing eccentricity. This space-variant resolution perception phenomenon - named foveation - is a hardwired mechanism and a natural way of reducing the amount of information streamed to the brain, in order to cope with neuronal transmission bandwidth limitations and the brain machinery processing capacity. However this compression phenomenon introduces a space variant uncertainty either in low-level visual processes as 3D reconstruction, or in higher-level object classification and recognition tasks. In order to efficiently explore and understand the surrounding environment [1], humans have developed a set of visual attention and oculo-motor mechanisms that allow them to actively direct the eyes towards different locations in the surrounding environment, thus cleverly compensating for the aforementioned sensory and computational limitations.

Likewise, robots deployed in everyday environments, are faced with increasingly complex scenarios where objects are arranged in many possible different spatial configurations. Moreover, the problem of deciding which regions in the visual

field are to be attended during the visual search task is computationally demanding. Therefore, like biological systems, robots should be endowed with a set of mechanisms that allow them to search for objects of interest and build detailed maps of the scene, while avoiding the potential computational overload of processing irrelevant sensory stimuli. Under the assumption that biological systems perform quasi-optimally in their environment due to multiple generations of genetic improvement, researchers have been developing biologically inspired systems [2] equipped with smart multi-resolution sensor topologies and provided with computational models of selective attention [3]. These implementations not only mimic the mechanisms observed in humans but also lead to more effective behaviours with constrained resources (computational, energetic), which is one of the main goals to pursue of today's cognitive robotics.

In this work we propose a probabilistic selective attentional framework for artificial systems provided with binocular foveal vision. Our framework relies on visual information and associated confidence measures (see Fig. 1) that are used to autonomously drive the agent's gaze direction during search tasks. Our contributions are the following. First, we model the stereo reconstruction uncertainty that arises as a result of spatial quantization phenomena inherent in the retina. Our approach considers Gaussian Receptive Fields¹ (RFs) distributed in space following two different tessellations: (i) a classical uniform (Cartesian) arrangement and (ii) a log-polar one that mimics the human retina. The RFs in the latter present a space-variant spatial distribution and support radius [5]. The Unscented Transform (UT) [6] is used to propagate belief from the 2D retina domain to 3D via stereo reconstruction. When compared with previous approaches that also assume Gaussian quantization noise and that rely on first order linearizations to approximate the non-linear transformations involved in 3D reconstruction [7], our method based on UT is more precise and hence improves 3D estimation quality. Second, the probabilistic sensory measurements are integrated in a novel versatile randomized Sensory Ego-Sphere (SES) whose topology can be biased according to the autonomous agent short-term tasks and goals. The proposed SES, helps achieving the task, by allocating the limited resources to

¹*Receptive fields* are the fundamental visual processing units. Each corresponds to a specific region in the retina (image) and is represented by the average value of the photo-receptors (pixels) within it (e.g. average color). For more details, we refer the interested reader to [4].

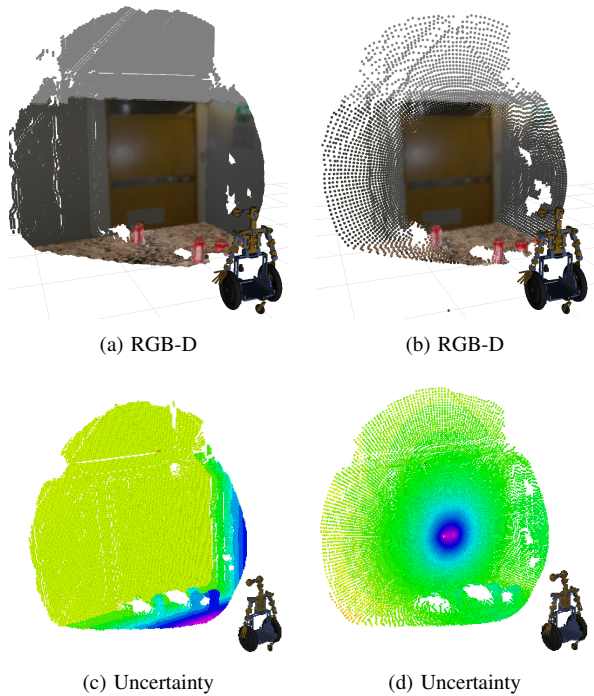


Fig. 1: A snapshot of the RGB-D point clouds (top row) and associated probabilistic measures (bottom row) obtained with the proposed Cartesian (left column) and foveal (right column) stereo sensor models. Blue and purple colors correspond to higher precision measurements.

important surrounding regions according to the task. Finally, a decision-making process, framed within a multi-armed bandit setting [8], acts as a mediating cognitive attentional process that seeks to maximize expected task-related rewards. The proposed decision algorithm relies on statistical measures to decide where to look next by selecting the most promising regions to attend. We investigate a simple Upper Confidence Bound (UCB) algorithm [9] for the task of finding the closest object to the observer. The UCB algorithm controls the exploration-exploitation trade-off typical of decision under uncertainty algorithms: to accomplish the task it is necessary to explore the world, but too much exploration will delay the task execution.

The remainder of this paper is organized as follows. In section II we conduct a brief overview of the attentional frameworks available in the literature with a strong emphasis on probabilistic-based methodologies. In section IV, we outline the proposed sensor observation model and the uncertainty propagation model from the retinal domain to 3D. In section V, we introduce a novel biologically inspired short-term memory structure which is egocentric, compact, and convenient for fast and efficient information update and retrieval. In section VI, we endow our system with a decision-making process that actively drives the agent’s gaze direction, through sequential saccadic eye movements. Finally, in section VII, we

experimentally validate our model and compare a conventional Cartesian camera against a space-variant vision system. The obtained results demonstrate that a wider field of view at the cost of less peripheral resolution is advantageous in visual-search tasks. We show that with our methodologies different gaze patterns emerge depending on the sensor characteristics and decisions on confidence bounds. Furthermore, we demonstrate that spatial memory biases, reflecting prior knowledge about the world structure and the task at hand, allow large performance improvements in visual search tasks.

II. RELATED WORK

Probabilistic based active vision requires not only the characterization of the sensory-motor uncertainties, but also the definition of memory structures that facilitate continuous recall and temporal fusion of probabilistic sensory data. Therefore, we organize the present section in two distinct parts. At first we overview the state-of-the-art in active vision with an emphasis on probabilistic models of overt attention. Afterward, we analyse the memory data structures proposed in the literature suitable for applications related to attention.

A. Overt Attention and Active Vision

The first studies on visual attention date back to the mid 19th century, pioneered by Hermann von Helmholtz [10] and motivated by the willingness to understand how humans attended stimuli at the periphery of the visual field. Nowadays, the literature on selective attention is vast, and covers a wide range of scientific fields, including cognitive neuroscience [11] and computer science [12], playing an important role in computer vision and robotics applications [13]. In this work we focus on a particular aspect of attention: active sensing which is tightly coupled to the concept of overt attention. The goal of active vision systems is to direct the eyes towards locations such that:

- the information about the surrounding environment is increased over time (exploration);
- the desired region is centered in the eyes and thus observed by the retinal zone of maximum visual acuity (exploitation).

Early work on computational modeling of overt attention [14] suggests that saccadic eye movements are attracted in a bottom-up manner to salient stimuli, i.e., to areas of the the visual field that differ from the neighborhood in different feature modalities and spatial scales, in a center surround fashion. Arguably, the most influential work on saliency modeling was proposed in [15] where the authors introduced a unique spatial saliency map which encoded the prominent locations over the entire visual scene, to be attended after and analyzed in detail in order of decreasing conspicuity. A major drawback of the saliency models is that the classical sequencing mechanism (inhibition of return) only considers spatial information (location) and not of the information provided by the stimuli. Therefore, they do not leverage information from previous steps to improve information on unexplored areas. Ongoing sensory data gathering while attending salient regions should continuously affect reasoning and influence decisions of where to look next.

This decision process requires not only a continuous evaluation of new stimuli arising from previously not visible regions, as well as reassessing previously attended locations as a result of the newly acquired information.

More sophisticated models are framed within probabilistic paradigms that account both for sensori-motor uncertainties as well as the world intrinsic stochasticity and unpredictability. A common idea behind these models is that statistical objectives are the fundamental driving elements behind visual attention. From a Bayesian standpoint, attention seeks to actively infer the future actions that maximize the expected information gain given the spatio-temporal context. Therefore informational gain is itself the inner goal behind attention [16].

The probabilistic-based saliency model proposed in [17] suggests that surprising events or stimuli attract attention. The Kullback-Leibler (KL) divergence between prior and posterior beliefs is by convention used as a measure of surprise. However, surprise models are purely exogenous by nature since they react to observed stimuli. Active vision models based on optimal stochastic control principles pose the action selection problem within Bayes risk minimization framework, and differ on the chosen policies. On one hand, infomax algorithms [18] seek to maximize the expected accumulated future informational gain in fixed time-horizon. On the other hand greedy MAP policies consider only a one-step look ahead time window [19] and self-knowledge about the retinal acuity map to decide the best location to attend. A recent work on active sensing accounted also for behavioral costs [20], such as the energy and temporal costs incurred in choosing a given motor action.

Despite the demonstrated applicability of the previously mentioned approaches on target search tasks in monocular images, there are no works studying depth cues inferred by stereo vision, and the influence of foveal vision in the search strategies on binocular setups. The stereo reconstruction problem using foveated images has been addressed in the literature, namely in [21], where the authors have shown that it is possible to compute dense disparity maps from log-polar images. Nevertheless, with foveal images, stereo matching accuracy degrades in the image periphery. This motivates the need for modeling depth uncertainty in stereo reconstruction, due to space-variant discretization in foveated images and use this uncertainty to decide where to look next. In this paper we analyze the ability of active foveal stereo systems to accurately map the environment and efficiently execute visual search tasks. Some visual tasks are more naturally represented in 3D, for instance the search for nearby objects, as illustrated in this work. Therefore, the main contribution of the paper is the formulation of visual search tasks in 3D and a the development of novel methods for uncertainty propagation and spatial representations required for this purpose. We show that adequate retinal topologies and 3D spatial representation play a role in the speed of execution and accuracy of localization of targets in 3D search tasks.

B. Spatial Memory Data Structures

Memory plays a key role and is a core component of any cognitive architecture. In real-time decision-making problems

involving perception, autonomous agents rely on memory structures to store and query continuously obtained information in a robust and efficient manner.

The Sensory Ego-Sphere (SES) is an egocentric, short-term memory structure [22] that is convenient for sensory data fusion and that has been extensively used in robotics applications involving attention [23], [24]. In the attention domain and from a practical point of view, egocentric spherical representations offer several advantages in applications involving humanoid robots when compared to typical Cartesian representations such as regular occupancy grids, point clouds or octrees [25]. Spherical representations based on egocentric polar coordinate systems, are typically more compact (low memory requirements) and avoid the requirement of computationally expensive ray-casting techniques when dealing with visibility issues.

Different representations and data structures for the SES have been proposed in the literature (see Fig. 4). Typically, 2D array type structures based on spherical coordinate systems are used to represent the spherical surface [23]. These can be accessed in $\mathcal{O}(1)$ time and thus are appropriate for real-time applications. Yet, they are non-isotropic and therefore data is not stored uniformly over the surface (i.e. the resolution is higher near the poles). On the other hand, the geodesic dome type data structure [22] is isotropic and therefore can better approximate 3D shape. However, indexing becomes less trivial and less efficient due to its non-regular topology. To tackle this issue [26] proposed a hierarchical geodesic structure that can significantly speed-up access times. In another work [27] the authors proposed an egocentric log-spherical grid named Bayesian Volumetric map occupancy spherical grid, that was proven suitable for probabilistic multi-modal sensor fusion.

Nevertheless, none of the previous mentioned structures are easily reconfigurable for the implementation of task-dependent cognitive biases that either enhance or impair the storage and recall of information in short-term spatial memory [28]. This fact motivates the need for more sophisticated, task-biased, versatile memory structures. In section V we propose a novel memory representation that tackles this problem. Typical tessellations of the sphere include quasi-uniform icosahedral tessellations, less uniform spherical polyhedra or non-uniform latitude/longitude grids. All these forms are highly regular and structured, which limits their flexibility to implement arbitrary shapes. The method proposed in this paper is based on projecting in the sphere randomly generated points according to a mixture of 3D gaussian distributed points with an arbitrary number of components, focal points (means) and dispersions (covariances). This generates an irregular grid but we can define more freely areas on the sphere with varying degrees of density and dispersion. Our sampling scheme is easy to implement and allows for the fast creation of task-biased sensory ego-spheres. As opposed to the previously proposed deterministic counterparts, our SES relies on an easy to implement random sampling scheme that allows for the fast creation and real-time access of arbitrary reconfigurable topologies.

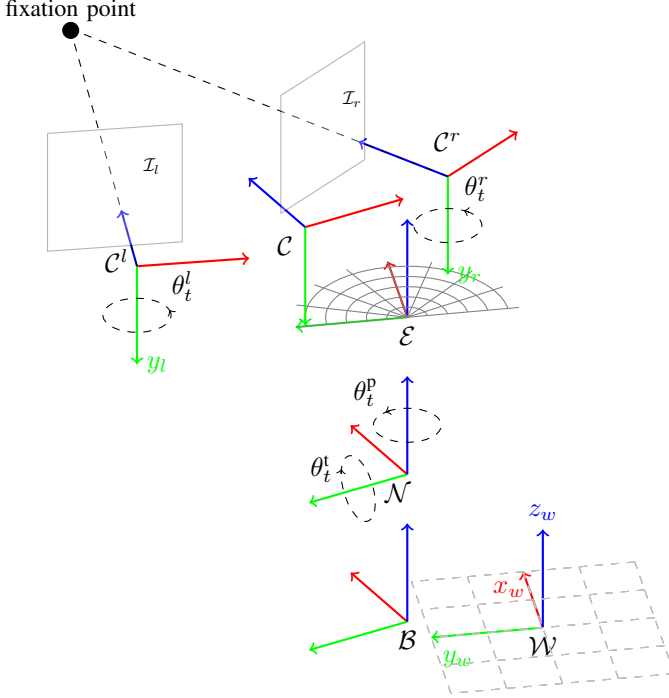


Fig. 2: The various coordinate systems used by our system (best seen in color): The inertial world frame (\mathcal{W}) in which the environment is represented; the base frame (\mathcal{B}) which is rigidly attached to the mobile robot base, and permits determining the robot pose in the world, given the odometric readings; the neck frame (\mathcal{N}) which allows representing pan and tilt cephalic movements; the egocentric frame (\mathcal{E}) in which spatial memory is embodied and sensor fusion is performed; the cyclopean frame (\mathcal{C}) in which stereo observations are represented; the convergent, non-parallel pair of cameras frames ($\mathcal{C}^l, \mathcal{C}^r$), in which monocular images are obtained.

III. PROBLEM STATEMENT AND SYSTEM OVERVIEW

In the proposed problem, the observer's goal is to select the oculomotor actions that maximize task related rewards. On one hand we rely on a recursive Bayesian filter that sequentially accumulates sensory inputs and extracts valuable information about the agent and the environment state, given noisy observations. On the other hand, a decision-making algorithm predicts the best future locations to gather information, according to some statistical or behavioral criteria.

The environment structure, i.e. 3D map, is a projection of the world structure $W \subset \mathbb{R}^3$ in the agent's egocentric reference frame \mathcal{E} , internally represented by a discrete set of points, each associated to a specific observation direction (see section IV). Let us denote the set of environment sample points by

$$X_t = \{\mathbf{x}_t^i \in \mathbb{R}^3, i = 1, \dots, N_x\} \quad (1)$$

where N_x is the total number of considered observation directions. These points are modeled as Gaussian random variables, initialized with mean and covariance selected according to

a priori knowledge about the type of environment in which the robot operates. The egocentric reference frame \mathcal{E} is head-centered, has three translational degrees of freedom and fixed orientation with respect to the world frame of reference (see Fig. 2).

In order to execute visual search tasks, the proposed cognitive architecture is equipped with two sensory-motor modalities:

- proprioception - provided by odometric and oculocephalic joint encoders;
- stereo vision - provided by a stereo camera system.

The observer is allowed to change its state, i.e. the observation view point, through base and oculocephalic movements (please see Fig. 1). At each time instant, the proprioceptive modality reports the robot base location and its internal kinematic state. More specifically, the robot position and orientation $P \in \mathbb{R}^6$ in the inertial frame of reference \mathcal{W} , the agent's eyes horizontal vergence ($\theta_t^v \in \mathbb{R}$) and the head pan and tilt joint angles ($\theta_t^p, \theta_t^l \in \mathbb{R}$). Let us denote the joint set of odometric and oculocephalic measured/controlled joint positions by

$$U_t = \{P_t, \theta_t^v, \theta_t^p, \theta_t^l\} \quad (2)$$

We assume that the proprioceptive modality provides noise-free observations. In other words, we consider that the measurement errors are negligible with respect to the visual sensor errors and therefore that the robot location and kinematics, and thus, the transformations between the various reference frames involved in our system (see Fig. 2), can be deterministically determined from U_t . Furthermore, we assume that the environment W is static for the duration of the search task and is not affected by the robot motor actions U_t (the base location and the posture of the robot's head). The preceding assumptions yield the following probabilistic simplification

$$p(X_t|W, U_t) = p({}^{\mathcal{E}}\mathbf{R}_{\mathcal{W}}W + {}^{\mathcal{E}}\mathbf{t}_{\mathcal{W}}) = p(X_t) \quad (3)$$

where ${}^{\mathcal{E}}\mathbf{R}_{\mathcal{W}} \in \mathbb{R}^{3 \times 3}$ and ${}^{\mathcal{E}}\mathbf{t}_{\mathcal{W}} \in \mathbb{R}^{3 \times 1}$ are an orthogonal rotation matrix and a translation vector, respectively, obtained by combining deterministic proprioceptive joint angle measurements with known forward kinematics.

The stereo sensor computes a list of 3D point estimates defined in a cyclopean reference frame \mathcal{C} , with origin at the midpoint of the stereo baseline, from noisy point correspondences observed in the left and right retinal domain. Let us denote the set of 3D points by

$$Z_t = \{\mathbf{z}_t^o \in \mathbb{R}^3, o = 1, \dots, N_{v,t}\} \quad (4)$$

where $N_{v,t}$ is the total number of observed independent and identically distributed (iid) measurements by the stereo sensor at time t . The observation model described in section V explains how measurements Z_t are generated according to the environment 3D structure, egocentric projection X_t :

$$Z_t \sim p(Z_t|X_t) \quad (5)$$

IV. STEREO SENSOR MODEL

In stereo vision, a general stereo matching algorithm computes a set of one-to-one point correspondences between two images [29]. However the precision of the measurements is finite and constrained by the fundamental image-sensing units size and spacing. In order to model reconstruction uncertainty due to the limited sensing precision at the retinal level we consider a probabilistic observation model for our stereoscopic sensor [30].

A. Nonparallel Stereo System

Let us suppose that our stereoscopic system is composed by a convergent, non-parallel pair of pinhole cameras $\mathcal{C}^l, \mathcal{C}^r$, allowed to rotate around their y optical-axis by $\theta^l = \frac{\theta^v}{2}$ and $\theta^r = -\frac{\theta^v}{2}$, respectively, and are separated by a fixed baseline b . Furthermore, let us assume that the stereo system is calibrated, and thus, the intrinsic $\mathbf{K}^l, \mathbf{K}^r$ and extrinsic $\mathbf{R}_y(\theta^v), \mathbf{T}(b)$, camera parameters are always known.

B. Gaussian Stereoscopic Retinal Observation Model

Let us consider that the cameras image planes $\mathcal{I}^l, \mathcal{I}^r \subset \mathbb{R}^2$ comprise a finite set of RFs denoted by $\mathcal{S}^l, \mathcal{S}^r \subset \mathbb{R}^2$. We assume that each RF has a non-uniform stimuli response, modeled by a two dimensional Gaussian profile [5], with the support regions depicted in Fig. 3. The mean $\boldsymbol{\mu} = (\mu_x, \mu_y)$ defines the coordinates of the center of the RF in the retinal plane, where response is maximal, and the standard deviation σ represents its support radius.

Thus, observing a correspondence at a given RF pair $\mathbf{s}^i \in \mathcal{S}^l \times \mathcal{S}^r$ follows a conditional Gaussian distribution:

$$\mathbf{s}^i \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{s}^i}, \boldsymbol{\Sigma}_{\mathbf{s}^i}) \quad (6)$$

where

$$\boldsymbol{\mu}_{\mathbf{s}^i} = \begin{bmatrix} \mu_{x,i}^l \\ \mu_{y,i}^l \\ \mu_{x,i}^r \\ \mu_{y,i}^r \end{bmatrix}, \quad \boldsymbol{\Sigma}_{\mathbf{s}^i} = \text{diag}(\sigma^{l,i^2}, \sigma^{l,i^2}, \sigma^{r,i^2}, \sigma^{r,i^2}) \quad (7)$$

C. Stereoscopic Reconstruction

Given a point pair correspondence \mathbf{s}^i found in the retinal domain, we determine the corresponding 3D position in the cyclopean reference frame via stereo analysis. Since point pair correspondences are inherently corrupted with precision errors, their projection lines may no longer satisfy the epipolar constraint and therefore not intersect in 3D space. Hence, one should rely on a triangulation method, denoted by τ , in order to compute a 3D Cartesian point estimate $\hat{\mathbf{z}}$ from a point correspondence in image coordinates \mathbf{s}^i :

$$\tau: \mathcal{I}^l \times \mathcal{I}^r \rightarrow \mathbb{R}^3 \quad (8)$$

Due to its simplicity and relatively low computational complexity, we use the mid-point method (for details please refer to [31]).

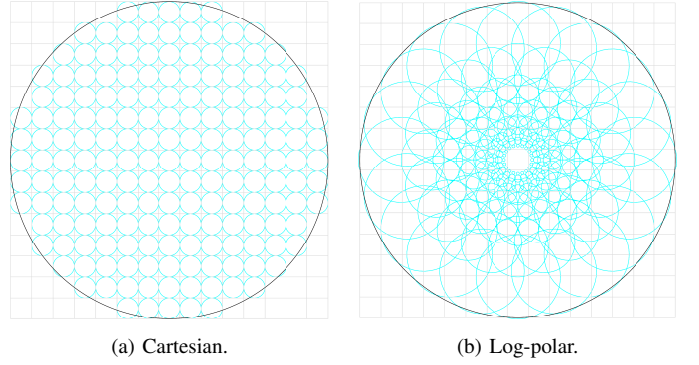


Fig. 3: Gaussian receptive fields with support plotted for 3 standard deviations

D. Uncertainty back-propagation via the Unscented Transform

Since the transformation (8) involved in 3D reconstruction is non-linear, we employ the Unscented transform [6] to compute the propagated mean and covariance up to the third order (by Taylor's expansion). This is achieved by approximating a multivariate Gaussian distributed variable with a set of meaningful and deterministically chosen set of samples (usually named sigma points). For each receptive field pair $\mathbf{s}^i \in \mathcal{S}^l \times \mathcal{S}^r$ we associate a set of sigma points

$$\mathcal{U}^i = \{\mathcal{X}^{(i,j)} \in \mathcal{I} \times \mathcal{I}^r : j = 0, \dots, 2N_s\} \quad (9)$$

where N_s is the number of sigma points, which are pre-computed according to the following expressions

$$\mathcal{X}^{(i,0)} = \boldsymbol{\mu}_{\mathbf{s}^i} \quad (10)$$

$$\mathcal{X}^{(i,j)} = \boldsymbol{\mu}_{\mathbf{s}^i} + \left(\sqrt{(N_s + \lambda) \boldsymbol{\Sigma}_{\mathbf{s}^i}} \right)_j \quad \text{for } j = 1, \dots, N_s \quad (11)$$

$$\mathcal{X}^{(i,j)} = \boldsymbol{\mu}_{\mathbf{s}^i} - \left(\sqrt{(N_s + \lambda) \boldsymbol{\Sigma}_{\mathbf{s}^i}} \right)_j \quad \text{for } j = N_s + 1, \dots, 2N_s \quad (12)$$

where $(\cdot)_j$ denotes the j -th row of a matrix. Furthermore, we consider a set of weights

$$\mathcal{W} = \{w_c^{(j)}, w_m^{(j)} \in \mathbb{R} : j = 0, \dots, 2N_s\} \quad (13)$$

which are computed as follows

$$w_m^{(0)} = \frac{\lambda}{L + \lambda} \quad (14)$$

$$w_c^{(0)} = \frac{\lambda}{L + \lambda} + (1 - \alpha^2 + \beta) \quad (15)$$

$$w_m^{(j)} = w_c^{(j)} = \frac{1}{2(L + \lambda)} \quad \text{for } j = 1, \dots, 2N_s \quad (16)$$

where $\lambda = \alpha^2(L + K) - L$ is a scaling factor, α controls the spread of the sigma points around the mean, K is a secondary scaling parameter, and β is used to incorporate prior knowledge about the distribution of \mathbf{s} (for Gaussian distributions $\beta = 2$ is optimal). Then, for a given point

correspondence in retinal domain, we first apply the non-linear transformation τ to the sigma points associated with the corresponding RF pair, $\mathcal{Z}^{(i,j)} = \tau(\mathcal{X}^{(i,j)})$, and then re-estimate the mean and covariance in the 3D domain, according to

$$\hat{\mu}_{\mathbf{z}^i} = \sum_{j=0}^{2N_s} w_m^{(j)} \mathcal{Z}^{(i,j)} \quad (17)$$

$$\hat{\Sigma}_{\mathbf{z}^i} = \sum_{j=0}^{2N_s} w_c^{(j)} \left(\mathcal{Z}^{(i,j)} - \hat{\mu}_{\mathbf{z}^i} \right) \left(\mathcal{Z}^{(i,j)} - \hat{\mu}_{\mathbf{z}^i} \right)^T \quad (18)$$

The sigma points in retinal domain are computed offline and stored in a linear array in order to speed up on-line uncertainty propagation.

V. RANDOMIZED SENSORY EGO-SPHERE

In the proposed framework the SES plays an intermediate role between the stereo sensor and the decision planning process. Probabilistic data arriving from the sensory stream is continuously fused and integrated over time in the SES, by means of recursive Bayesian filter. At the same time the available information is used to predict and redirect gaze to the best expected location, in the light of new observations.

A. Definition

The proposed SES is composed of a set of cells \mathcal{P} lying on a unit sphere, whose center corresponds to a certain absolute orientation, and a map that assigns to each cell the 3D coordinates of the point observed by the robot at that orientation

$$\mathcal{M} : \mathcal{P} \longrightarrow \mathcal{R}^3 \quad (19)$$

The proposed cell grid structure is analogous to a Voronoi diagram defined on a spherical 2-manifold \mathbb{S}^2 in 3D space, as depicted in Fig. 4. In practice the proposed SES comprises a set of 3D Cartesian sample points with unit norm and centered in the observer egocentric reference frame \mathcal{E} ,

$$\mathcal{P} = \{ \mathbf{p}^i \in \mathbb{R}^3, i, \dots, N_x : \|\mathbf{p}^i\| = 1 \} \quad (20)$$

which are i.i.d. and randomly generated from a three dimensional Gaussian Mixture Model (GMM) distribution

$$\mathbf{p}^i = \frac{\mathbf{v}^i}{\|\mathbf{v}^i\|} \text{ where } \mathbf{v}^i \sim p(\theta) = \sum_{m=1}^M \phi^m \mathcal{N}(\mu_p^m, \Sigma_p^m) \quad (21)$$

where M is the number of mixture components and where each $\mathbf{p}^i \in \mathcal{P}$ represents an orientation, allowing for efficient data-alignment with observed 3D points, using inner products (equation 23). Each SES cell, represented by $\mathbf{p}^i \in \mathcal{P}$, stores one environment sample point estimate $\mathbf{x}^i \in X$.

The statistics of the Gaussian Mixture Model distribution are chosen according to the observer goals. On one hand, in order to produce uniform and unbiased memory structures, the surface should be sampled from a rotationally symmetric distribution, i.e., from a single Gaussian with zero mean

and variance equal in all dimensions [32] (Fig. 4d). On the other hand, non-uniform, task-dependent memory biasing can be achieved by manipulating the Gaussian Mixture Model parameters. The proposed randomized representation offers a convenient mechanism for encoding task and world prior knowledge. Memory biasing should lead to more efficient, flexible and adaptable memory allocation and to more effective behaviours during task execution.

Hypothetical topologies that may be suitable for different tasks are depicted in Fig. 4: If for instance the task is to look for people, one should privilege areas at the equator rather than the poles. In this case, varying the Gaussian mean is not sufficient. One could sample from a single-component zero mean GMM with larger variance in the horizontal directions (Fig. 4e). While crossing a street, the observer should prioritize attentional resources to antipodal, lateral regions (Fig. 4f). This can be achieved by sampling from a single-component Gaussian with a larger variance in the lateral component, or from a two-component GMM with opposite lateral means. More complex tasks can benefit from irregular topologies with multiple foci, obtained from GMMs with many components (Fig. 4g).

B. Data Alignment

For each observed world point estimate provided by our stereo observation model at time t , \mathbf{z}_t^o , we need to find the associated memory cell in order to perform probabilistic data fusion. The association process goes as follows. First, the observed random variable is transformed from the cyclopean to the egocentric reference frame, according to the linear transformation $Z' : \mathbb{R}^3 \longrightarrow \mathbb{R}^3$ of the form

$$Z' = {}^{\mathcal{E}}\mathbf{R}_C Z + {}^{\mathcal{E}}\mathbf{t}_C \quad (22)$$

where ${}^{\mathcal{E}}\mathbf{R}_C \in \mathbb{R}^{3 \times 3}$ is an orthogonal rotation matrix and ${}^{\mathcal{E}}\mathbf{t}_C \in \mathbb{R}^{3 \times 1}$ a translation vector, obtained by combining proprioceptive joint angle measurements with known forward kinematics.

Second, for each observation \mathbf{z}_t^o we find the associated memory cell c^o , which is the one that minimizes the Euclidean distance, according to the mapping function \mathcal{M} , here defined as follows

$$c^o = \mathcal{M}(\hat{\mu}_{\mathbf{z}_t^o}) = \underset{j}{\operatorname{argmin}} \langle \mathbf{p}^j, \frac{\hat{\mu}_{\mathbf{z}_t^o}}{\|\hat{\mu}_{\mathbf{z}_t^o}\|} \rangle \quad (23)$$

After finding the associated cell we update its respective estimate according to equation (33). Moreover, we assume that the transformed observations are conditionally independent, given X_t , and thus

$$p(Z'_t | X_t) = \prod_{o=1}^{N_v} p(\mathbf{z}_t^o | \mathbf{x}_t^{c^o}) \quad (24)$$

Finally, the resulting probabilistic observation model $p(\mathbf{z}_t^o | \mathbf{x}_t^{c^o})$ follows a Gaussian distribution

$$\mathbf{z}_t^o | \mathbf{x}_t^{c^o} \sim \mathcal{N}(\hat{\mu}_{\mathbf{z}_t^o}, \hat{\Sigma}_{\mathbf{z}_t^o}) \quad (25)$$

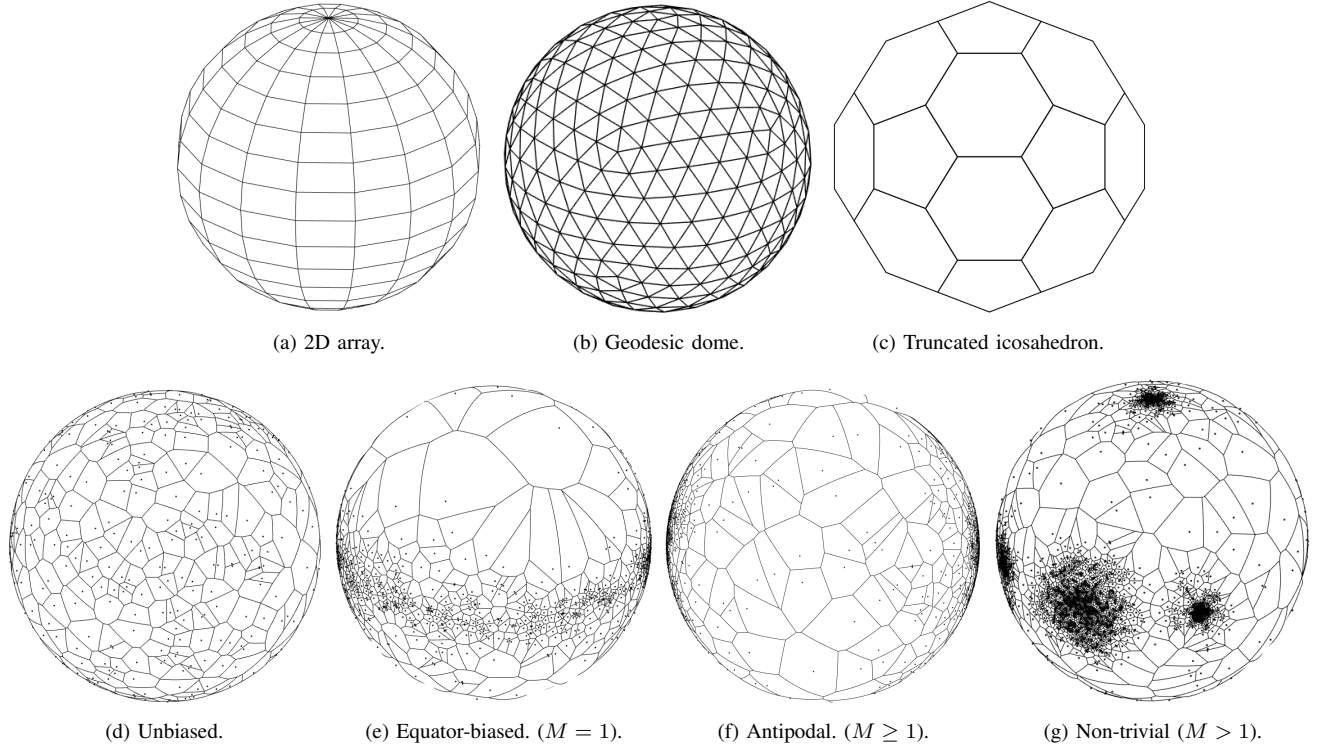


Fig. 4: Different Sensory Ego-Spheres, resulting from different tessellations: top row illustrates highly regular, deterministic structures. The bottom row depicts our novel randomized structure for different task-dependent biases.

with statistics computed as follows

$$\hat{\mu}_{\mathbf{z}'_t} = \varepsilon \mathbf{R}_C \hat{\mu}_{\mathbf{z}_t} + \varepsilon \mathbf{t}_C \quad (26)$$

$$\hat{\Sigma}_{\mathbf{z}'_t} = \varepsilon \mathbf{R}_C \hat{\Sigma}_{\mathbf{z}_t} \mathbf{R}_C^T \quad (27)$$

C. Probabilistic Sensor Fusion

In the sensor fusion perspective, the goal of the optimal Bayesian estimator is to determine the posterior probability distribution over X , given the accumulated visual sensory observations and the robot proprioceptive state measurements up to time $t \in \mathbb{N}$. Sequential Bayesian filtering allows us to accumulate sensor inputs and update the likelihood of X , at each time instant.

The posterior probability distribution at time t , of the set of internal environment sample points X_t given the current and past visual and proprioceptive observations, is given by

$$p(X_t | Z_{1:t}, U_{1:t}) = p(X_t | Z_t, Z_{1:t-1}, U_{1:t}) \quad (28)$$

Furthermore, since we assume that the proprioceptive measurements are deterministic, then

$$p(Z'_t | Z_t, U_t) = p(\varepsilon \mathbf{R}_C Z_t + \varepsilon \mathbf{t}_C) = p(Z'_t) \quad (29)$$

and equation (28) becomes

$$p(X_t | Z'_{1:t}) = p(X_t | Z'_t, Z'_{1:t-1}) \quad (30)$$

Since the world is static, at each iteration, the solution to the filter involves only one update step: in the *measurement update* step observations are used to update the current belief by applying the Bayes rule to the right hand side of equation (30) and using the observation model (5) we get

$$p(X_t | Z'_{1:t}) = \eta p(Z_t | X, Z'_{1:t-1}) p(X | Z'_{1:t-1}) \quad (31)$$

where η is a normalizing constant. Since the current observations Z'_t are conditionally independent of the past observations $Z'_{1:t-1}$ given the current environment projection in the egocentric frame, X_t , the previous equation becomes

$$p(X_t | Z'_{1:t}) = \eta p(Z'_t | X_t) p(X_t | Z'_{1:t-1}) \quad (32)$$

The *a posteriori* is independently determined for each cell, according to

$$p(\mathbf{x}_t^{c^o} | \mathbf{z}'_{1:t}) = \eta p(\mathbf{z}'_t | \mathbf{x}_t^{c^o}) p(\mathbf{x}_t^{c^o} | \mathbf{z}'_{1:t-1}) \quad (33)$$

and follows a Gaussian distribution, with statistics given by

$$\hat{\Sigma}_{\mathbf{x}_t^{c^o}} = \left(\hat{\Sigma}_{\mathbf{x}_{t-1}^{c^o}}^{-1} + \hat{\Sigma}_{\mathbf{z}'_t}^{-1} \right)^{-1} \quad (34)$$

$$\hat{\mu}_{\mathbf{x}_t^{c^o}} = \hat{\Sigma}_{\mathbf{x}_t^{c^o}} \left(\hat{\Sigma}_{\mathbf{x}_{t-1}^{c^o}}^{-1} \hat{\mu}_{\mathbf{x}_{t-1}^{c^o}} + \hat{\Sigma}_{\mathbf{z}'_t}^{-1} \hat{\mu}_{\mathbf{z}'_t} \right) \quad (35)$$

Each point estimate is initialized with large mean and covariance to reflect the high uncertainty due to non-existent world prior knowledge.

VI. ACTIVE SENSING: SEQUENTIAL STOCHASTIC DECISION MAKING

In the proposed framework, the decision making is responsible for sensory-motor coordination. Based on probabilistic information stored in memory, the decision process selects where to look next and the associated desired motor commands. Like other approaches that use uncertainty and task-related rewards to guide decision making we frame our approach within the reinforcement learning domain [33]. As such, the agent selects the action that maximizes expected task related cumulative rewards.

The underlying framework for decision making under uncertainty, which assumes non-deterministic noisy state observations, is known as *Partially Observable Markov Decision Process* (POMDP). In our particular problem formulation we have continuous states and observations, as well as large discrete action spaces (i.e. each SES cell represents an action), which renders intractable the computation of optimal policies. Even approximate methods for solving POMDPs would take considerable time (e.g. hours or days). Moreover, if the environment changes the observer needs to recompute the full policy. Hence, state-of-the-art methods for solving POMDPs are unsuitable for large problems, which require real-time on-line decision making. Therefore, rather than framing our problem as a POMDP, we rely on simpler and less costly tools from Bayesian Optimization, for reinforcement learning. More concretely, from multi-armed bandit problems (MAB) [34].

A. Saccadic Planning as a Multi-armed bandit Problem

In MAB problems, at each time instant the agent selects an action and collects a reward. The rewards are drawn from a posterior probability distribution whose statistics are continuously updated over time. Typically, the goal of the agent is to maximize the sum of collected rewards or, equivalently, minimize cumulative regret.

In this work the selected task was to find the closest object to the observer as fast (i.e. with minimum fixations) and precisely (i.e. with minimum uncertainty) as possible. Within the MAB framework, this is commonly referred to as the best-arm identification problem [35].

In our particular setting, each world sample point represented in memory is a bandit whose statistics are not known in advance. The agent chooses actions, i.e. a fixation point, from the set of alternatives $a \in \{1, \dots, N_x\}$ and collects payoffs from a reward distribution $r(\mathbf{x}^a)$. Considering the task at hand, we define the reward obtained when choosing a given action a as a function of the distance to the ego-frame

$$r(\mathbf{x}^a) = -\|\mathbf{x}^a\|_2 \quad (36)$$

Since \mathbf{x}^a follows a Gaussian distribution, then we consider a first order approximation for the reward distribution such that

$$r(\mathbf{x}^a) \sim \mathcal{N}(\mu_r(\mathbf{x}^a), \sigma_r(\mathbf{x}^a)) \quad (37)$$

where $\mu_r(\mathbf{x}^a)$ and $\sigma_r(\mathbf{x}^a)$ are computed as follows

$$\mu_r(\mathbf{x}^a) = E[-\|\mathbf{x}^a\|_2] = -\|\hat{\mu}_{\mathbf{x}^a}\|_2 \quad (38)$$

$$\sigma_r(\mathbf{x}^a) = \text{Var}[-\|\mathbf{x}^a\|_2] \approx \mathbf{J}^T \hat{\Sigma}_{\mathbf{x}^a} \mathbf{J} \quad (39)$$

where $E[\cdot]$ and $\text{Var}[\cdot]$ denote the expectation and variance operators, respectively, and \mathbf{J} is a Jacobian matrix, defined as follows

$$\mathbf{J} = \left. \frac{\partial r(\mathbf{x})}{\partial \mathbf{x}} \right|_{\hat{\mu}_{\mathbf{x}^a}} = \left[\frac{x_{\hat{\mu}_{\mathbf{x}^a}}}{\|\hat{\mu}_{\mathbf{x}^a}\|_2} \frac{y_{\hat{\mu}_{\mathbf{x}^a}}}{\|\hat{\mu}_{\mathbf{x}^a}\|_2} \frac{z_{\hat{\mu}_{\mathbf{x}^a}}}{\|\hat{\mu}_{\mathbf{x}^a}\|_2} \right]^T \quad (40)$$

B. Acquisition Functions

In the Bayesian optimization framework, acquisition functions are responsible for defining the strategy when searching for the optimum. The literature on acquisition functions used to guide stochastic optimization is vast and includes many different heuristics that deal with the exploration-exploitation dilemma. On one hand, Probability of Improvement (PI) [36] methods select the action that maximizes the probability of improving the current instantaneous reward. On the other hand, Expected Improvement (EI) [37] seeks for the action that maximizes the expected improvement magnitude. More recently, the idea of using Upper Confidence Bounds (UCB) [38] to deal with exploration-exploitation trade-offs in machine learning problems has proven successful in robotics applications [39], exhibiting increased preference for exploration when compared to the former approaches. Since the best performing acquisition function is highly dependent on the objective at hand, the authors in [40] propose combining single acquisition functions in mixed portfolio strategies.

In this work we compared three different action selection strategies:

- 1) a simple yet powerful UCB algorithm named "Sequential Design for Optimization" [41] that is easy to implement and elegantly handles the trade-off between exploration (minimizing uncertainty) and exploitation (maximizing rewards) that emerges in decision making under uncertainty [9]. At each time instant, the observer selects the alternative with maximal upper confidence bound on the expected reward, given the past observations, according to the following expression

$$a_t = \underset{a \in \{1, \dots, N_x\}}{\text{argmax}} \mu_r(\mathbf{x}_t^a) + \alpha \sigma_r(\mathbf{x}_t^a) \quad (41)$$

where α is a user selected parameter that controls the width of the confidence bound and thus the exploration behaviour during task execution.

- 2) The probability of improvement, which at each time instant, selects the action with highest probability of leading to an improvement upon the current best (\mathbf{x}_t^*), as follows

$$a_t = \underset{a \in \{1, \dots, N_x\}}{\text{argmax}} \mathbb{P}(r(\mathbf{x}_t^a) > r(\mathbf{x}_t^*)) \quad (42)$$

- 3) The expected improvement, which tries to maximize the expected magnitude of the improvement upon the so far

best, according to

$$a_t = \underset{a \in \{1, \dots, N_x\}}{\operatorname{argmax}} \mathbb{E}(r(\mathbf{x}_t^a) - r(\mathbf{x}_t^*)) \quad (43)$$

$$(44)$$

Finally, the motor-action U_t corresponding to $\mathbf{x}_t^{a_t}$ is computed from known forward kinematics.

VII. RESULTS

In order to demonstrate the applicability of the proposed framework and compare the performance of different visual sensor topologies, we performed a set of experiments in simulation. In all of the experiments we constrained the number of RFs - and hence the computational resources - to be always fixed and equal in the Cartesian and log-polar cases (please refer to [5] for mathematical details on the log-polar distribution). We considered $N_{rf} = 200 \times 200$ images in both cases.

The remainder of this section is organized as follows. We begin by characterizing and assessing the ability of the different sensors to map the environment with low uncertainty. Then, we proceed to evaluating the performance of the complete active task-oriented stereo sensing framework, in a realistic simulated environment.

A. Sensor Characterization

To characterize the proposed sensor model, we assessed the average uncertainty in 3D reconstruction as a function of depth, vergence angle and sensor type in the following manner: First, we generated a set of fronto-parallel planar surfaces, with varying distance $d \in [0, 1]$ from the binocular system. Depth is constant for all points lying within the same planar-surface.. Then, for each planar surface we varied the vergence angle, in the interval $\theta^v \in [0, \frac{\pi}{2}]$ and computed the corresponding 3D reconstructions with associated uncertainties. Note that in this experiment we are not characterizing a full environment but just single snapshots taken by the observer. Furthermore, we assumed that an object can be approximated by a planar surface occupying the observing agent field of view. This allows for comparing both sensors, under the same conditions.

Let us consider the log-determinant of the inverse covariance matrix (also known as precision matrix) to quantify pointwise information:

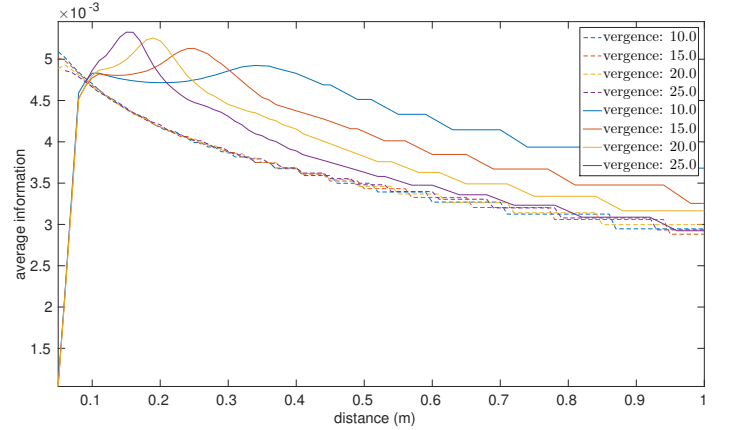
$$I(\Sigma) = -\log(|\Sigma|) \quad (45)$$

Here we rely on the average information gathered with a single depth image to assess the quality of the sensors, which is defined as follows

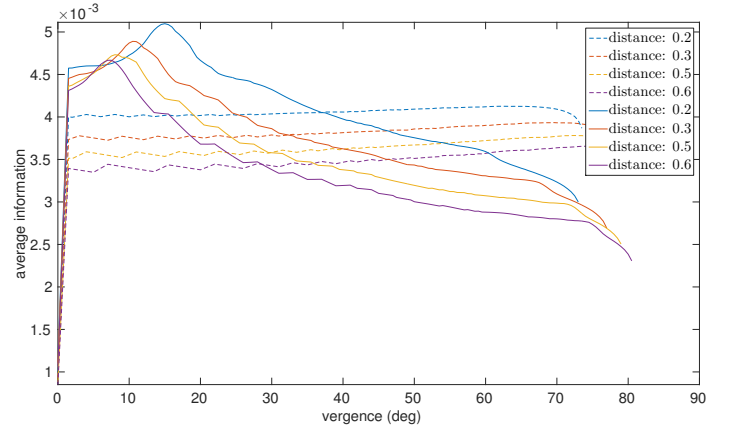
$$TI = \frac{1}{N_{rf}} \sum_{i=1}^{N_{rf}} I(\hat{\Sigma}^i) \quad (46)$$

As depicted in Fig 5, the foveal outperforms the Cartesian sensor, in terms of gathered information which is maximal if the fixation point coincides with the planar surface. Furthermore, the Cartesian sensor information reliability does not

depend on the vergence angle and decays monotonically with increasing depth, while the foveal sensor performance critically depends on the careful selection of the vergence angle. These results are directly in line and support previous findings [42] that suggest that foveal distributions facilitate stereo vision in convergent systems. In foveated systems gaze acts like a focus of attention, which when directed to the point of interest, improves dramatically the depth resolution around the fixation point. Instead, optical vergence movements in Cartesian systems provide no gains in 3D resolution, resulting only in unnecessary energetic costs.



(a) Varying distance for different vergence angles curves.



(b) Varying vergence for different planar distance curves.

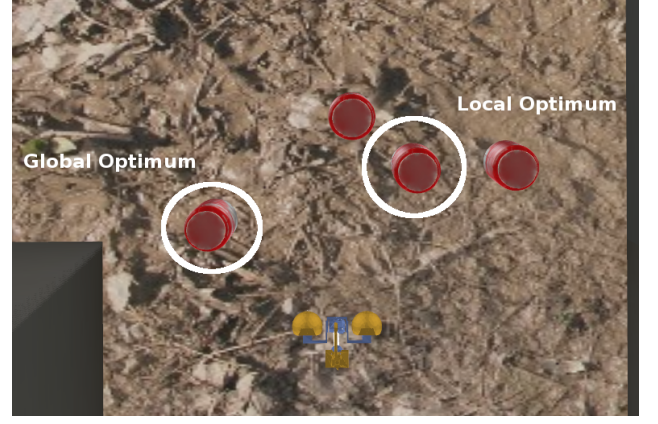
Fig. 5: Numerical characterization of the sensor model for the Cartesian (dashed lines) and the Log-polar (solid lines) sensors, as a function of distance and vergence.

B. Active Vision

With the view of investigating how our methodology performs in simultaneous target searching and mapping, we performed a set of experiments in the Gazebo simulator with the Vizzy robot [43] head (see Fig. 6). For the sake of the experiments simplicity, the robot was fixed to the ground



(a) The simulation scenario created for evaluating the proposed active vision framework.



(b) The global optimum was placed at a non-trivial location which could only be attended with either sufficient exploration or a wide field of view.

Fig. 6: The simulation scenario created for evaluating the proposed active vision framework. The task to perform was to find the nearest object from the robot ego frame. The evaluation scenario contained a non-trivial global optimum which could only be attended if enough exploration was promoted.

floor and hence the motion was restricted to oculocentric movements. However, note that our methodology is also applicable to scenarios in which the robot platform can move. This would imply updating the 3D point estimates stored in memory taking into account the uncertainty in robot base movements (odometry), and implementing a z-buffer technique to determine which point to store in each cell, due to possible occlusions occurring after translations.

We created a static scenario with multiple objects (coke cans) displaced at arbitrary depths, over a highly textured background, in order to facilitate stereo reconstruction, which is highly dependent on the environment texture richness. The Gazebo simulator generates pinhole camera images, with uniform resolution. Hence, for the log-polar sensor, we generated foveated images from uniform resolution images by first applying the log-polar transformation and then converting back to Cartesian domain via the inverse transformation. This operation has the effect of blurring the image in the periphery while maintaining high resolution in the center. Finally, disparity maps were computed using a state-of-the-art dense stereo matching algorithm named Semi-Global Block Matching (SGBM) [44].

As previously pointed out, the task at hand was to find the nearest world point to the observer. Points on the ground floor are easily excluded by thresholding the z_w coordinate. In all experiments we fixed the number of memory sample points to $N_x = 20000$. In each experiment we let the observer perform $T = 50$ saccadic movements, with initial ($t = 1$) pan, tilt and vergence angles equal to zero. Each experiment was repeated 20 times in order to average out variability in different real-time simulations. Non-repeatability was influenced by multiple factors including separate threads for Gazebo's physics and sensor generation, as well as stochastic delays involved in higher level inter-process communication. Furthermore, in

order to deal with motion blur and visuo-proprioceptive delays that arise during saccadic eye movements, we used the visual suppression mechanism proposed in [45], which temporarily blinds the observer during saccades.

1) *Evaluation Metrics:* In order to quantitatively assess the performance of our methodologies we considered the following evaluation metrics:

- the gap reduction metric [46] which is a quality measure that evaluates how effectively the algorithm is at finding the global maximum:

$$g_t = \frac{\mu_r(\mathbf{x}^+) - \mu_r(\hat{\mathbf{x}}_1^{a_1})}{\mu_r(\mathbf{x}^*) - \mu_r(\hat{\mathbf{x}}_1^{a_1})} \quad (47)$$

where $\mu_r(\mathbf{x}^*)$ is the true global maximum

$$\mu_r(\mathbf{x}^*) = \max_i \mu_r(\mathbf{x}^i) \quad (48)$$

and $\mu_r(\mathbf{x}^+)$ is the best obtained reward up to time t

$$\mu_r(\mathbf{x}^+) = \max_t \mu_r(\hat{\mathbf{x}}_t^{a_t}) \quad (49)$$

The gap is defined between 0, meaning no improvement over the initial fixation, and 1 for the optimal improvement. In order to measure the speed for task completion and thus performance efficiency we also assess the average gap reduction per saccade which implicitly represents the average progress towards the optimum per saccade:

$$G/S = \frac{1}{T} \sum_{t=1}^T g_t \quad (50)$$

- the cumulative regret which is a standard metric, here suitable to evaluate the convergence behaviour during the search for the optimum:

$$R_t = \mu_r(\mathbf{x}^*) - \frac{1}{t} \sum_{k=1}^t \mu_r(\hat{\mathbf{x}}_k^{a_k}) \quad (51)$$

Notice that here we are not interested in minimizing the total regret, i.e. the incurred losses during exploration, but instead on finding the global optimum. When normalized by the number of saccades it represents the temporal cumulative regret gain per saccade:

$$R/S = \frac{1}{T} \sum_{t=1}^T R_t \quad (52)$$

- the average global gathered information which is a quality performance measure of the global knowledge gathered about the world up to time t (exploratory behaviour):

$$GI_t = \frac{1}{N_x} \sum_{i=1}^{N_x} I(\hat{\Sigma}_t^i) \quad (53)$$

When normalized by the number of saccades it represents the temporal average global information gain per saccade:

$$GI/S = \frac{1}{T} \sum_{t=1}^T GI_t \quad (54)$$

- the nearest object gathered information, which is a target reconstruction quality measure that benefits high precision (i.e. low uncertainty) in target reconstruction:

$$LI_t = I(\hat{\Sigma}_t^i) \quad \forall i: \|\hat{\mu}_{\mathbf{x}_t^*} - \hat{\mu}_{\mathbf{x}_t^i}\| < R_{NN} \quad (55)$$

where $\hat{\Sigma}_t^*$ is the true known global maximum estimated covariance at time t and R_{NN} is a user-selected nearest-neighbor radius. We considered $R_{NN} = 0.1m$ in all the experiments described below.

When normalized by the number of saccades it represents the temporal average local information gain per saccade:

$$LI/S = \frac{1}{T} \sum_{t=1}^T LI_t \quad (56)$$

2) *Foveal vs Cartesian*: Our first aim was to compare the behaviour of the foveal against the Cartesian sensor during task execution, for different upper confidence bound parameter values $\alpha \in \{0, 0.01, 1, 100, \infty\}$ and different sensing field of views $\text{fov} \in \{90^\circ, 135^\circ\}$. The sensor field of views were selected such that in one of the cases ($\text{fov} = 90^\circ$) the global optimum was not in the field of view of the observer at $t = 1$. The SES cells were generated from a unbiased, zero mean Gaussian distribution at initialization (see Table I).

A global analysis of the results depicted in Fig. 8 shows that the foveal sensor outperforms the Cartesian both in terms of the quality of the gathered information, as well as the task execution speed and effectiveness, as demonstrated by the gap

reduction plot. We hypothesize that the best performance of the foveal sensor is due to the fact that the uncertainty in the periphery implicitly promotes more peripheral (lateral) exploration whereas the Cartesian promotes longitudinal (depth) search. This statement is clearly supported by the cumulative regret plots which exhibit lower losses for the Cartesian sensor, and thus a greedier behaviour. Moreover, for the foveal sensor case, a larger FOV allows the agent to attend the target more quickly at the cost of reduced information gain. A wider FOV, despite having less peripheral resolution, is advantageous in the speed of execution during visual search tasks.

In Fig. 8a we assess the performance of our method for the different acquisition functions referred in section VI. On one hand, in the UCB case, a larger confidence bound parameter α increases exploration and, on average, improves performance in the particular task of finding the nearest object. However, too much exploration incurs in large cumulative regrets, and thus in high energy costs due to large oculoccephalic movements when attending objects further from the observer. Nevertheless, purely exploratory behaviours ($\alpha = \infty$) lead to better results in the average reconstruction quality as shown by the information metrics, since on average more memory sample locations are fixated. On the other hand, the tested improvement-based policies (PI and EI) seek to improve on the current best and have the advantage of being parameter free. For our particular setting, and similarly to UCB with $\alpha = 0$, PI tends to be excessively greedy and get trapped in local minima. On the contrary, EI deals well with the exploration-exploitation trade-off, as demonstrated by the average gap reduction and cumulative regret per saccade metrics due to the fact that it implicitly accounts for the improvement magnitude of each saccadic action, which allows for choosing distant, with high variance, fixation points.

An in-depth analysis of the temporal evolution metrics (Fig. 8b) for a fixed $\alpha = 100$, allows us to assess convergence times for a fairly exploratory behaviour. The temporal evolution of the gap reduction metric shows that, in all cases, no more than 20 saccades are necessary to perform the task of finding the nearest object for both sensor types. Howbeit, as indicated by the accumulated regret temporal evolution, convergence is only achieved after no less than 30 saccades. We further note that, after convergence, the cumulative regret is on average higher for the foveal case, as a consequence of having a more exploratory nature. The gathered information exhibits an asymptotically convergence behaviour and has a faster transient time for the Cartesian sensor, again, supporting the idea that the Cartesian sensor is more greedy, myopic, and thus more prone to get trapped in local minima.

3) *Memory Biases*: Here our goal was to investigate the effect of different spatial memory topological biases imposed from *a priori* knowledge regarding the environment structure and the task at hand. At the present, experiments were performed with the foveal sensor with a $\text{fov} = 135^\circ$, and for the UCB with $\alpha = 100$. We intended to demonstrate that a careful displacement of the memory patches considering prior knowledge about the surrounding environment and the task at hand should incur in large performance gains. Therefore, we considered four different prior belief distributions with

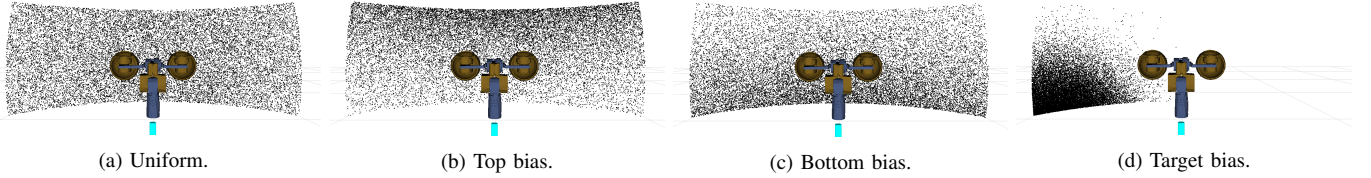


Fig. 7: SES sample point distribution according to different topological memory biases and kinematic constraints.

Bias	μ_p			Σ_p		
	x	y	z	xx	yy	zz
unbiased	0	0	0	0.5	0.5	0.5
top	0	0	1	0.5	0.5	0.5
down	0	0	-1	0.5	0.5	0.5
target	0.61	0.43	-0.67	0.05	0.05	0.05

TABLE I: Memory biasing parameters.

parameters defined in Table I and resulting SES topologies depicted in Fig. 7:

- a “neutral” (unbiased) distribution reflecting the absence of a priori knowledge about the target location.
- a “bad” (top) prior belief distribution based on the wrong assumption that the object is above the observer.
- a “good” (down) prior belief distribution that assumes that the object is on the ground
- a “very good” (target) prior belief considering the true location of the target object.

In the Fig. 9 we can observe that the “target” case had the best performance and the “top” the worst performance according to all metrics. In fact, as demonstrated by the gap reduction and the accumulated regret time evolution plots, the method was successful in finding the global optimum and converged with only 2 saccades. All the other cases were still able to find the optimum with less than 10 saccades and converge to the optimum within the first 20 saccades.

As expected, the gathered average local information metric indicates that increasing the memory sample density around the object of interest improves the target’s gathered information. These experiments demonstrate that translating task-related priors in clever memory allocation to regions of higher reward yields faster task execution times and faster convergence rates. This results in an increase in the time spent on reducing the uncertainty on the target and therefore in improved reconstruction quality. On the one hand, promoting higher resolution in spatial memory to the most important surrounding regions according to the task, allows for more accurate target reconstruction. On the other hand, less fixations are needed to find the target, since less memory cells, and thus possible fixations, will reside outside of the target vicinity.

VIII. CONCLUSION

In this work we investigated the impact of uncertainty due to quantization phenomena in the retina and on how to take

advantage of it to guide gaze shifts for two distinct retinal topologies: Cartesian and log-polar. With our approaches different gaze patterns emerge depending on the sensor topology and field of view and on exploration-exploitation confidence bounds parameters. The obtained results demonstrate that a wider field of view, despite less peripheral resolution is advantageous in visual search tasks execution speed. Furthermore, we showed that a task-biased SES allows for simultaneously coping with limited memory resources (i.e. limited number of memory cells) while improving performance, both in terms of target reconstruction quality and task execution speed.

The proposed framework can be further enhanced with other ideas from the attentional stereopsis literature, namely with [47] which improves reconstruction quality and efficiency by restricting stereo matching to biologically plausible volumes of interest. Moreover, one could improve the proposed SES run-time performance by relying on a nearest neighbour data alignment scheme. A kd-tree could be built during initialization time and be used for storing the ego-sphere cells (\mathcal{P}). Then, for each observed 3D point, searching for the closest cell on the sphere would be performed in $O(\log N_x)$, instead of $O(N_x)$. Future work will include developing adaptive re-sampling techniques for on-line memory biasing, capable of coping with dynamic tasks and environments.

ACKNOWLEDGMENT

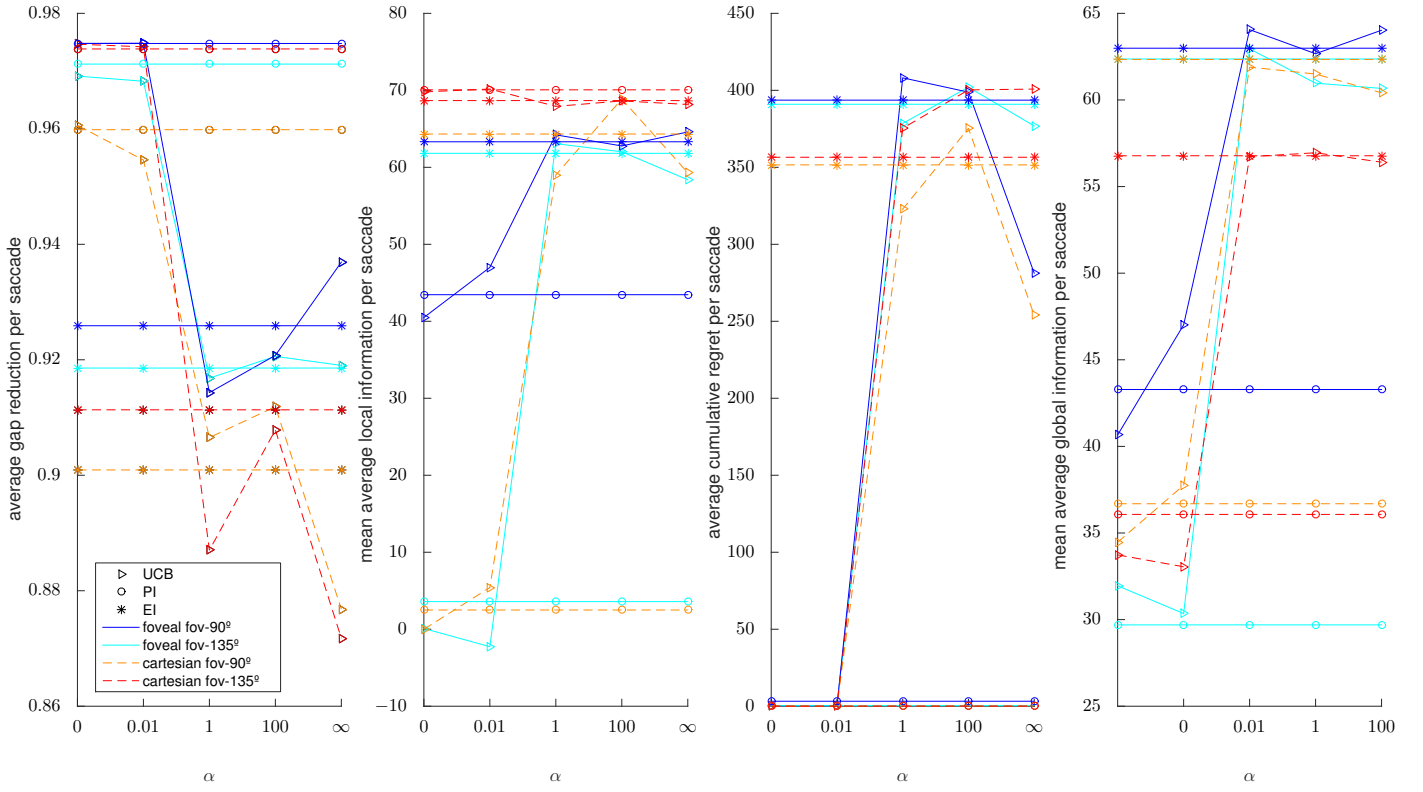
This work has been partially supported by the Portuguese Foundation for Science and Technology (FCT) project [UID/EEA/50009/2013]. Rui Figueiredo is funded by FCT PhD grant PD/BD/105779/2014. Helder Araújo would like to thank FCT (Portuguese Foundation for Science and Technology) grant UID-EEA- 0048-2013.

REFERENCES

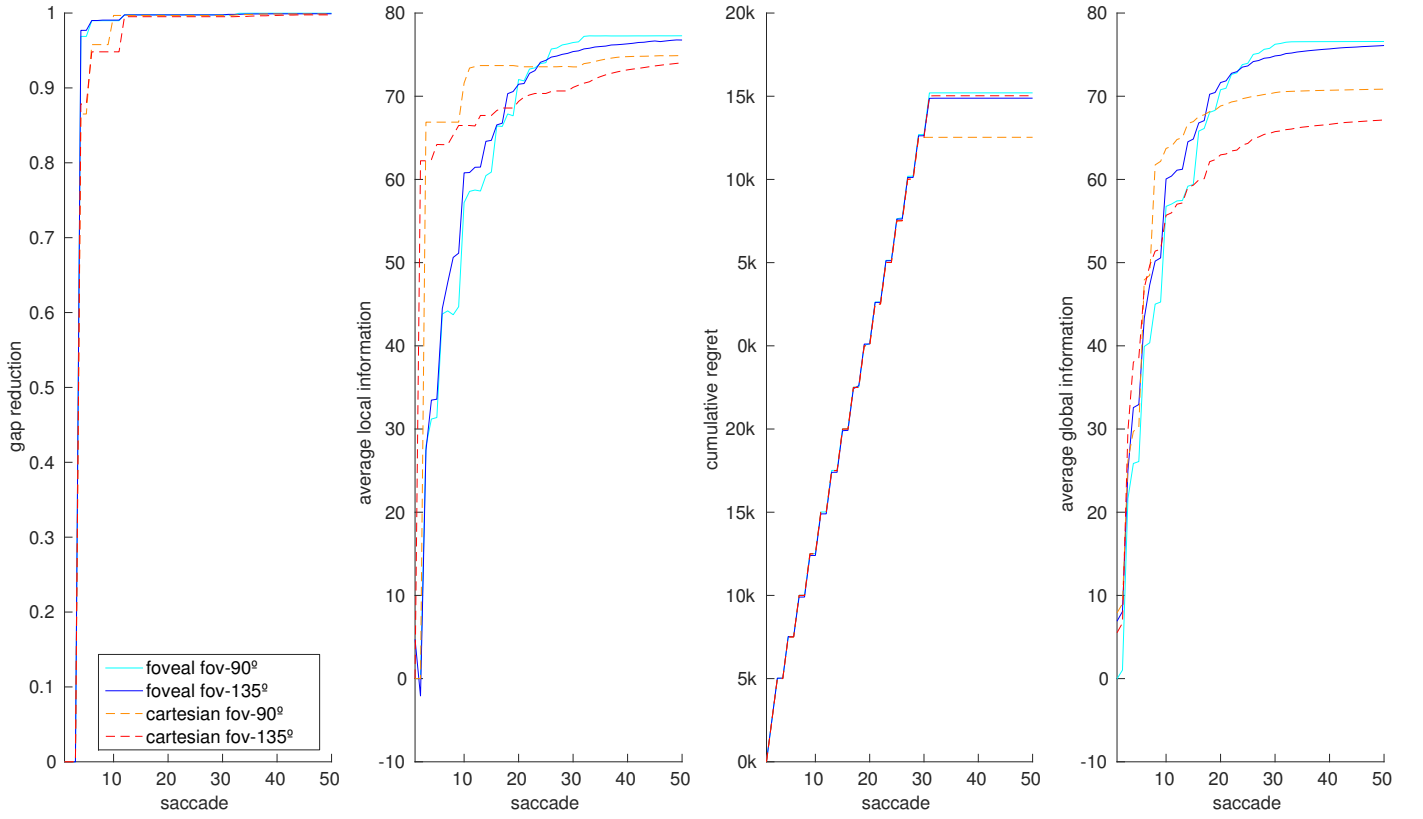
- [1] M. Posner, *Cognitive Neuroscience of Attention*. Guilford Press, 2012.
- [2] C. Colombo, M. Rucci, and P. Dario, “Integrating selective attention and space-variant sensing in machine vision,” in *Image Technology*, Springer, 1996, pp. 109–127.
- [3] S. Vijayakumar, J. Conradt, T. Shibata, and S. Schaal, “Overt visual attention for a humanoid robot,” in *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, IEEE, vol. 4, 2001, pp. 2332–2337.

- [4] S. Edelman, "Receptive fields for vision: From hyper-acuity to object recognition," 1995.
- [5] D. Pamplona and A. Bernardino, "Smooth foveal vision with gaussian receptive fields," in *9th IEEE-RAS International Conference on Humanoid Robots, Humanoids 2009, Paris, France, December 7-10, 2009*, 2009, pp. 223–229.
- [6] S. Julier and J. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, Mar. 2004.
- [7] D. J. Kriegman, E. Triendl, and T. O. Binford, "Stereo vision and navigation in buildings for mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 6, pp. 792–803, 1989.
- [8] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.
- [9] R. Agrawal, "Sample mean based index policies with $o(\log n)$ regret for the multi-armed bandit problem," *Advances in Applied Probability*, pp. 1054–1078, 1995.
- [10] H. Von Helmholtz, *Handbuch der physiologischen Optik*. 1866, vol. 9.
- [11] M. Carrasco, "Visual attention: The past 25 years," *Vision Research*, vol. 51, no. 13, pp. 1484–1525, 2011, Vision Research 50th Anniversary Issue: Part 2.
- [12] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 1, pp. 185–207, Jan. 2013.
- [13] M. Begum and F. Karray, "Visual attention for robotic cognition: A survey," *Autonomous Mental Development, IEEE Transactions on*, vol. 3, no. 1, pp. 92–105, Mar. 2011.
- [14] C. Koch and S. Ullman, "Shifts in selective visual attention: Towards the underlying neural circuitry," in *Matters of intelligence*, Springer, 1987, pp. 115–141.
- [15] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 11, pp. 1254–1259, 1998.
- [16] K. Friston, R. Adams, and R. Montague, "What is value – accumulated reward or evidence?" *Frontiers in Neurobotics*, vol. 6, no. 11, 2012.
- [17] L. Itti and P. F. Baldi, "Bayesian surprise attracts human attention," in *Advances in Neural Information Processing Systems, Vol. 19 (NIPS*2006)*, Cambridge, MA: MIT Press, 2006, pp. 547–554.
- [18] N. J. Butko and J. R. Movellan, "Infomax control of eye movements," *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 2, pp. 91–107, 2010.
- [19] J. Najemnik and W. S. Geisler, "Optimal eye movement strategies in visual search," *Nature*, vol. 434, no. 7031, pp. 387–391, 2005.
- [20] S. Ahmad and A. J. Yu, "Active sensing as bayes-optimal sequential decision making," *CoRR*, vol. abs/1305.6650, 2013.
- [21] A. Bernardino and J. Santos-Victor, "A binocular stereo algorithm for log-polar foveated systems," English, in *Biologically Motivated Computer Vision*, ser. Lecture Notes in Computer Science, H. Bülthoff, C. Wallraven, S.-W. Lee, and T. Poggio, Eds., vol. 2525, Springer Berlin Heidelberg, 2002, pp. 127–136.
- [22] R. A. Peters li, K. A. Hambuchen, and R. E. Bodenheimer, "The sensory ego-sphere: A mediating interface between sensors and cognition," *Auton. Robots*, vol. 26, no. 1, pp. 1–19, Jan. 2009.
- [23] J. Ruesch, M. Lopes, A. Bernardino, J. Hornstein, J. Santos-Victor, and R. Pfeifer, "Multimodal saliency-based bottom-up attention a framework for the humanoid robot icub," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, May 2008, pp. 962–967.
- [24] K. A. Fleming, R. A. P. II, and R. E. Bodenheimer, "Image mapping and visual attention on a sensory ego-sphere," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2006, October 9-15, 2006, Beijing, China*, 2006, pp. 241–246.
- [25] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013, Software available at <http://octomap.github.com>.
- [26] M. Hirose, H. Furuhashi, T. Miyasaka, and K. Araki, "Reconstruction of range data by means of geodesic dome type data structure," *The Journal of the Institute of Image Electronics Engineers of Japan*, vol. 31, no. 3, pp. 388–395, 2002.
- [27] J. Ferreira, P. Bessière, K. Mekhnacha, J. Lobo, J. Dias, and C. Laugier, "Bayesian models for multimodal perception of 3d structure and motion," in *International Conference on Cognitive Systems (CogSys 2008)*, Karlsruhe, Germany, 2008.
- [28] L. E. Crawford, D. Landy, and A. N. Presson, "Bias in spatial memory: Prototypes or relational categories," *Poster presented at the 36th Annual Conference of the Cognitive Science Society, Quebec*, 2014.
- [29] B. Tippetts, D. J. Lee, K. Lillywhite, and J. Archibald, "Review of stereo vision algorithms and their suitability for resource-limited systems," *Journal of Real-Time Image Processing*, vol. 11, no. 1, pp. 5–25, 2016.
- [30] M. Perrollaz, A. Spalanzani, and D. Aubert, "Probabilistic representation of the uncertainty of stereo-vision and application to obstacle detection," in *Intelligent Vehicles Symposium (IV), 2010 IEEE*, Jun. 2010, pp. 313–318.
- [31] J. Wang and Y. Liu, "A closed-form solution of reconstruction from nonparallel stereo geometry used in image guided system for surgery," English, in *Multimedia Content Analysis and Mining*, ser. Lecture Notes in Computer Science, N. Sebe, Y. Liu, Y. Zhuang, and T. Huang, Eds., vol. 4577, Springer Berlin Heidelberg, 2007, pp. 371–380.
- [32] M. E. Muller, "A note on a method for generating points uniformly on n-dimensional spheres," *Commun. ACM*, vol. 2, no. 4, pp. 19–20, Apr. 1959.

- [33] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st. Cambridge, MA, USA: MIT Press, 1998.
- [34] H. Robbins *et al.*, "Some aspects of the sequential design of experiments," *Bulletin of the American Mathematical Society*, vol. 58, no. 5, pp. 527–535, 1952.
- [35] J.-Y. Audibert and S. Bubeck, "Best arm identification in multi-armed bandits," in *COLT-23th Conference on Learning Theory-2010*, 2010, 13–p.
- [36] H. J. Kushner, "A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise," *Journal of Fluids Engineering*, vol. 86, no. 1, pp. 97–106, 1964.
- [37] J. Mockus, "On bayesian methods for seeking the extremum," in *Proceedings of the IFIP Technical Conference*, London, UK, UK: Springer-Verlag, 1974, pp. 400–404.
- [38] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [39] D. Lizotte, T. Wang, M. Bowling, and D. Schuurmans, "Automatic gait optimization with gaussian process regression," in *In Proc. of IJCAI*, 2007, pp. 944–949.
- [40] M. D. Hoffman, E. Brochu, and N. de Freitas, "Portfolio allocation for bayesian optimization.," Citeseer.
- [41] D. D. Cox and S. John, "Sdo: A statistical method for global optimization," in *In Multidisciplinary Design Optimization: State-of-the-Art*, 1997, pp. 315–329.
- [42] C. F. R. Weiman, "Binocular stereo via log-polar retinas," SPIE, Ed., 1995.
- [43] P. Moreno, R. Nunes, R. Figueiredo, R. Ferreira, A. Bernardino, J. Santos-Victor, R. Beira, L. Vargas, D. Aragão, and M. Aragão, "Vizzy: A humanoid on wheels for assistive robotics," in *Robot 2015: Second Iberian Robotics Conference*, Springer International Publishing, 2016, pp. 17–28.
- [44] H. Hirschmuller, "Stereo processing by semiglobal matching and mutual information," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 328–341, 2008.
- [45] J. Avelino, R. Figueiredo, P. Moreno, and A. Bernardino, "On the perceptual advantages of visual suppression mechanisms for dynamic robot systems," in *International Conference on Biologically Inspired Cognitive Architectures (BICA)*, 2016.
- [46] D. Huang, T. T. Allen, W. I. Notz, and N. Zeng, "Global optimization of stochastic black-box systems via sequential kriging meta-models," *Journal of global optimization*, vol. 34, no. 3, pp. 441–466, 2006.
- [47] A. Agarwal and A. Blake, "Dense stereo matching over the panum band," *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, no. 3, pp. 416–430, 2010.

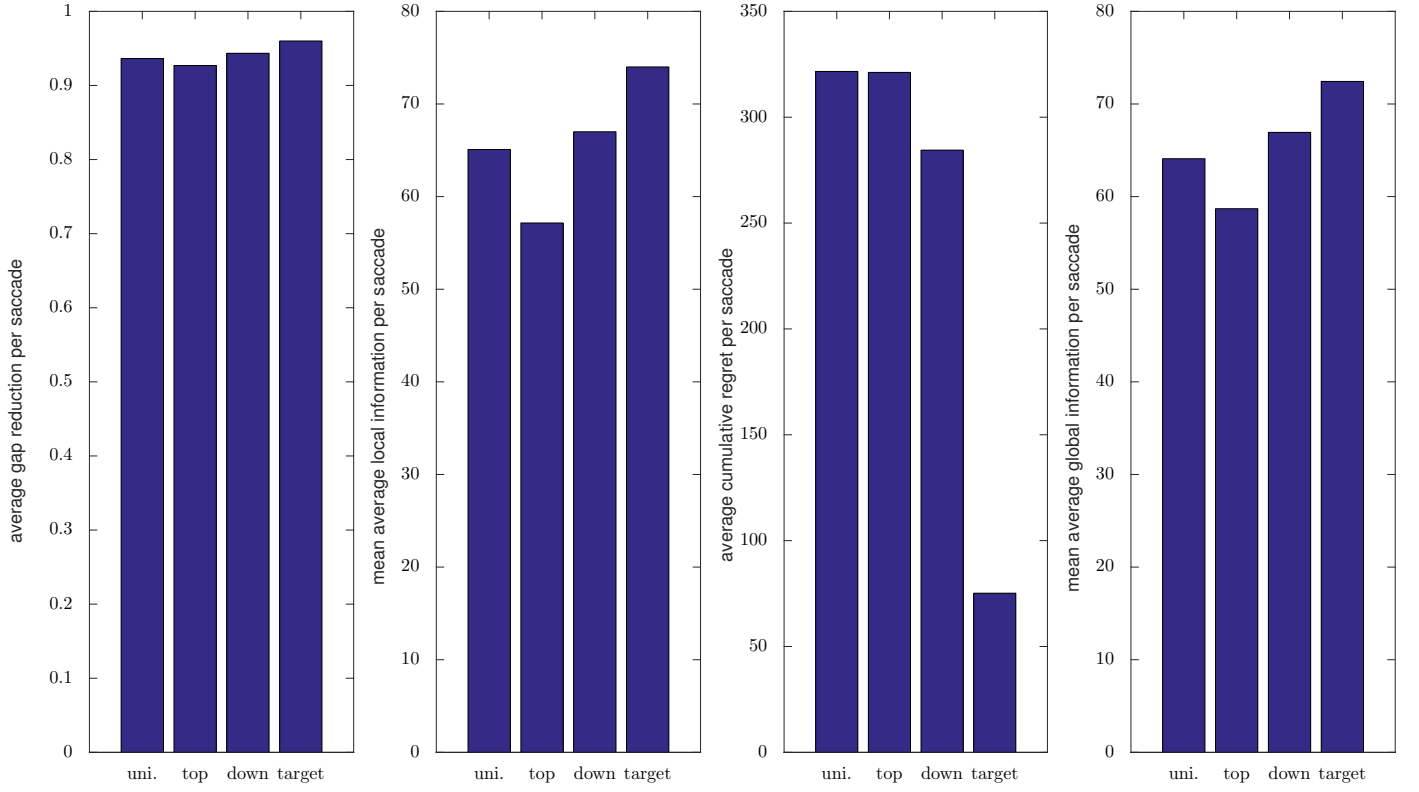
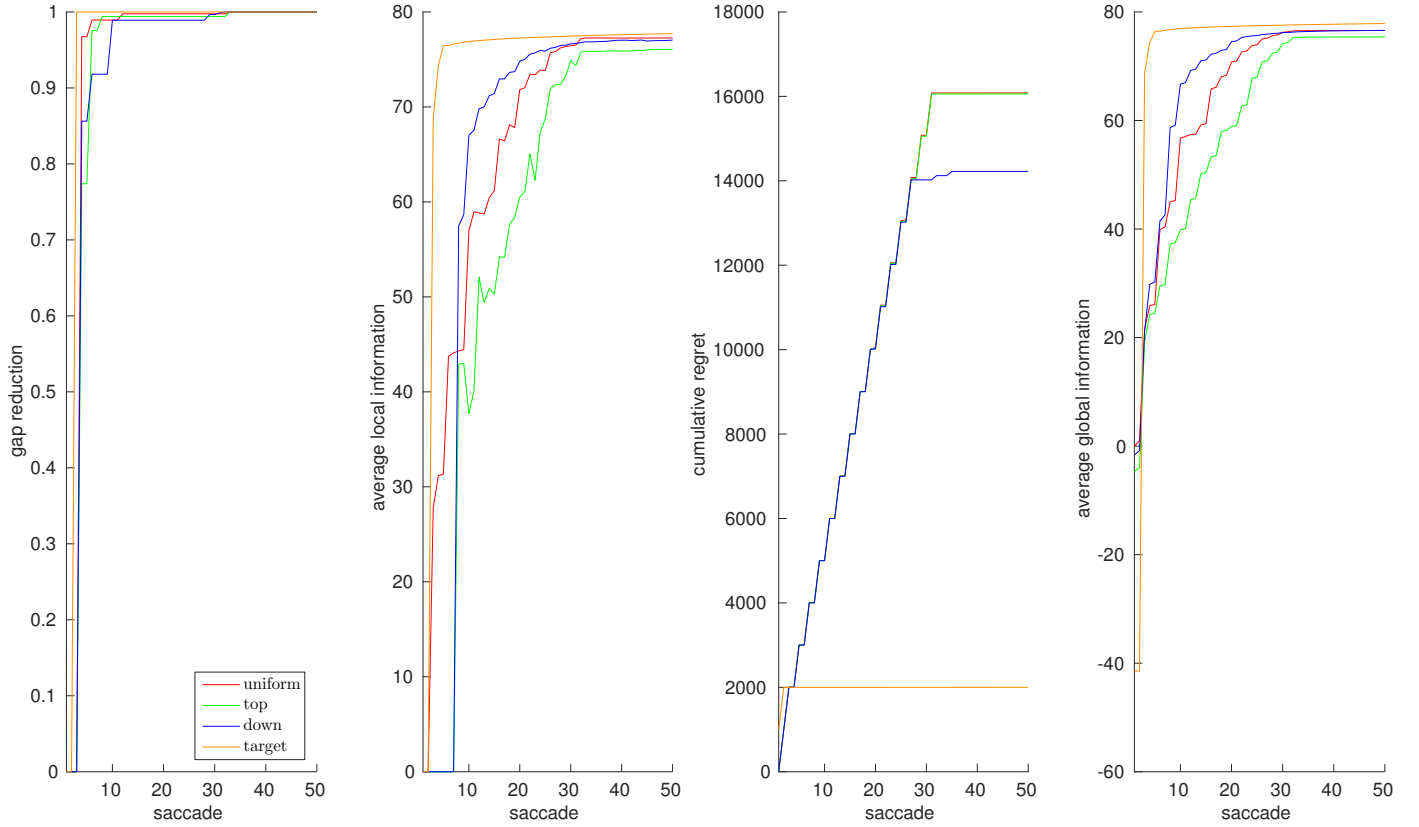


(a) Average per saccade performance plots.



(b) Time evolution plots.

Fig. 8: Performance results for the assessed sensor topologies, field of views and upper confidence bound parameter.

(a) Average per saccade performance plots. $\alpha = 100$ 

(b) Time evolution plots.

Fig. 9: Performance results for the assessed memory biases.