Spatial distribution of eye-movements after central vision loss is consistent with an optimal visual search strategy

October 9, 2019.

A. Vasilyev and M. Hansard

School of EECS, Queen Mary University of London,
London E1 4NS, UK.

The problem of gaze allocation has previously been studied in the framework of eye-movement control models, which require prior knowledge of visibility maps (VMs). These encode the signal-to-noise ratio, at each point in the visual field, which can be used to define an optimal policy of gaze allocation. However, it is not always possible to estimate the VM, in a given experimental setting, as it depends on many factors, including the visual system of the individual observer. Hence, few eye-movement datasets include the corresponding VM estimates. This can be problematic for the analysis of certain clinical conditions, such as Age-related Macular Degeneration (AMD), which are associated with reduced sensitivity in the affected locations of the visual field. The corresponding VMs are highly idiosyncratic, and cannot be modelled by estimates obtained from healthy observers. We propose an algorithm for maximum likelihood VM estimation, working directly from eye-movement sequences. We apply this algorithm to two eye-tracking datasets, based on visual search tasks, obtained from AMD patients. We show that the inferred VMs are spatially consistent with the measured visual field sensitivities. We also show that simulations with the estimated VMs can account for the asymmetric distribution of saccade vectors, which is typical of AMD patients.

Keywords: Oculomotor control, macular degeneration, visual search, Markov decision processes.

1. Introduction

Age-related macular degeneration (AMD) is a medical condition which may result in reduced acuity in areas of the visual field known as scotomas. As a result, the visual behaviour of AMD patients differs significantly from that of normal participants, during visual search tasks[1–3] and visual exploration of scenes.[4] AMD patients have lower performance in execution of visual tasks,[5] anomalous spatial distributions of saccade vectors[2] and larger numbers of saccades per trial.[4] Despite the considerable number of eye-tracking experiments conducted with AMD patients, there have been no attempts to explain the statistics of their saccades using computational models of eye-movements.[6, 7]

Control models of eye-movements[8, 9] aim to evaluate the optimal policy of gaze allocation, for a computational agent which has visual capabilities that are based on those of human observers. These models can produce eye-movements that are statistically consistent with human data, including response times on visual search tasks, spatial fixation distributions,[10] length and direction distributions of saccades,[11, 12] and geometrical and statistical persistence of eye-movements.[13] The control models of eye-movements are formulated as Partially Observable Markov Decision Processes (Po-Mdps), meaning that the agent has incomplete information about the underlying state of the visual environment. From the point of view of control models, visual search is a recurrent process of extraction and integration of visual information at each fixation, followed by a decision about where next to fixate. The main characteristic that governs the decision-making process of a computational agent in a Po-Mdp is the Visibility Map (VM). This is the ratio of perceptual signal-to-noise, across the visual field. Previously, it was shown that introduction of gaze-contingent simulated scotoma

(gaze-locked local blurring) results in adaptations of observers' eye-movement strategies[14, 15] which are consistent with the change of VM, due to the simulated scotoma. The use of control models of eye-movements implies that VMs are known for each experimental setting and participant. However, it is not always possible to estimate the required VMs, because they depend on the properties of the stimulus, the visual task, and the observer's visual system. For this reason, typical eye-movement datasets do not include corresponding VMs. Furthermore, in the case of clinical studies, the individual VMs may be highly idiosyncratic, owing to the diverse patterns of retinal degeneration.

In this paper we present a novel scheme for inference of individual VMs from sequences of eye-movements, which is inspired by the Policy Gradient with Parameter Exploration algorithm used for estimation of optimal policy in Po-Mdp.[16–19] This scheme estimates the parameters of the VM that maximize the probability of the observed eye-movement trajectories. We test this algorithm on both simulated and actual eye-movement data.[2, 3] In particular, we use an Entropy Limit Minimization ('Infomax') principle[11] as the basis for a generative model of eye-movements. We hypothesize that if the main difference between AMD patients and healthy controls is lower visual sensitivity across the visual fields of the former, then the difference in their visual behavior is due to their reduced ability to extract information about the visual scene. On the basis of analysis of eye-movements, Van der Stigchel et al. hypothesized that there is a preference for saccades towards parts of the visual field that are masked by scotomas, at a given moment.[2] This hypothesis is realistic in the Infomax scenario for the following reason. At each step of Po-Mdp, the probability distribution is updated according to the visual input; however, because the observer is effectively blind in scotomas, these areas cannot be eliminated as potential target locations. It follows that they may be preferred for the next fixation. However, using simulations of the Infomax model, we found that preference towards scotomas takes place only for the most trivial configurations of the VM. At the same time, we found that there is no evidence for a correlation between the angular distribution of eye-movements and the distribution of visual field values for AMD patients.

The present work is related to the Medusa algorithm for learning the observation parameters of a Po-Mdp.[20–22] This is based on matching the actions of an oracle (i.e. a system whose parameters are unknown) and a computational agent. Usage of this method in the case of stochastic execution of actions will result in noisy estimation of the objective function, with an error that grows with a dimensionality of action space. Our approach is based on the estimation of the probability of action sequence as a state space integral, in the Po-Mdp, unlike sampling methods in which actions of an oracle are compared to a given number of generated trajectories. We have developed an algorithm capable of inferring visibility maps in Po-Mdp with heuristic policies. We plan to extend this approach to the inference in Po-Mdp with optimal policies estimated with reinforcement learning methods.[8, 13, 23] We provide a full mathematical rationale of our method in supplementary material A.3, as well as the associated code.[24]

The goal of this research is to show that for each set of eye-movement trajectories, from a given human AMD patient, there exists a choice of visibility map that satisfies the following criteria:

- The simulation of an Infomax model with this VM explains the spatial distribution of eye-movements.
- The probability of an eye-movement sequence, estimated with this visibility map, exceeds chance level both in training and validation sets.
- This VM is consistent with measured visual field values, if available. In particular, areas with low signal-to-noise ratios correspond to scotomas, which can be identified using standard visual field tests.[25]

The basic statistical properties of AMD eye-movements have been comprehensively reported in previous works.[4, 5, 26–30] In particular, we use data provided by Janssen & Verghese[3] and Van der Stigchel et al.[2]

2.   Simulation of eye-movements

We use an Infomax model of eye movements,[11] in which the stochastic execution of saccades is driven by a greedy heuristic policy of information acquisition.[13]

2.1.   *Infomax model*

At the beginning of each episode (sequence of states and actions of an agent that ends in some terminal state), the target is placed at image location $\mathbf{u}_\star = (x_\star, y_\star)$, which is randomly picked from L possible locations:

$$\mathbf{u}_\star \in \{\mathbf{u}_1, \ldots, \mathbf{u}_\ell, \ldots, \mathbf{u}_L\}. \tag{1}$$

We assume that the target is placed on background noise or surrounded by distractors, placed at vacant locations. Each episode starts with a random initial location of gaze $\mathbf{x}_0 = (x_0, y_0)$, and finishes as soon as the fixation $\mathbf{x}_k$ corresponds to the target location $\mathbf{u}_\star$. Note that fixations can occur anywhere in the image, at discrete times t_k.

At each fixation $\mathbf{x}_k$, the observer receives sensory observations $S(\mathbf{x}_k, \mathbf{u}_\ell)$, with $\ell = 1, \ldots, L$, which can be abbreviated to

$$s_{k,\ell} = S(\mathbf{x}_k, \mathbf{u}_\ell). \tag{2}$$

These sensory observations are independent random variables, which represent the perceptual evidence that the target is at each of the L locations, on the k-th fixation. In particular, their values are distributed as[9]

$$p(s_{k,\ell} \mid \mathbf{x}_k) \sim N\left(\delta_{\mathbf{u}_\ell, \mathbf{u}_\star}, 1/V(\mathbf{u}_\ell - \mathbf{x}_k)\right) \tag{3}$$

where $N(\mu, \sigma)$ represents the 1D Normal distribution. The mean of the distribution is zero, unless the target is at location $\mathbf{u}_\ell$. This dependency is controlled by the Kronecker delta $\delta_{\mathbf{u}_\ell, \mathbf{u}_\star}$. The function $V(\mathbf{u} - \mathbf{x}) > 0$ is the visibility map (VM), which mediates the retinal information. The visibility, which is also known as the Fovea-Peripheral Operating Characteristic (FPOC),[8] is determined by the signal to noise ratio at location $\mathbf{u}$, given fixation at $\mathbf{x}$, as argued by Najemnik and Geisler.[11] Hence $1/V(\mathbf{u}_\ell - \mathbf{x}_k)$ gives the 'precision' of the sensory observation $s_{k,\ell}$. In normal observers, the visibility function is approximately radial, $V(\mathbf{u} - \mathbf{x}) = V_r(|\mathbf{u} - \mathbf{x}|)$, and is determined by retinal resolution and neural noise.

The decision making of the observer is modeled as a PO-MDP, with belief state $\mathbf{b}_k$ regarding the target location, at step k:

$$\mathbf{b}_k = (b_{k,1}, \ldots, b_{k,\ell}, \ldots, b_{k,L}). \tag{4}$$

This is a discrete probability distribution function, defined over image locations, given all observations received up to the k-th fixation. The probability distribution function is updated using Bayesian inference,[23] such that

$$b_{k,\ell} \propto p(s_{k,\ell} \mid \mathbf{x}_k)\, b_{k-1,\ell}. \tag{5}$$

The decision of where to fixate next is made on each step of PO-MDP, according to a policy of gaze allocation π:

$$\overline{\mathbf{x}}_{k+1} \leftarrow \pi(\mathbf{b}_k). \tag{6}$$

After making the decision, the actual coordinates of the next fixation location $\mathbf{x}_{k+1}$ are defined by execution function $\alpha(\mathbf{x})$:

$$\mathbf{x}_{k+1} = \alpha(\overline{\mathbf{x}}_{k+1}) = \overline{\mathbf{x}}_{k+1} + \boldsymbol{\xi}_{k+1} \tag{7}$$

where $\boldsymbol{\xi}_k$ is Gaussian-distributed spatial error, with zero mean and standard deviation ν. The standard deviation can be modelled as an affine function[31] of the intended saccade amplitude

$$\nu = \zeta_0 + \zeta_1 \|\mathbf{x}_{k+1} - \mathbf{x}_k\| \tag{8}$$

with parameters $\zeta_0 = 0.87°$ and $\zeta_1 = 0.084$ as estimated by Engbert et al.[31] The next step of PO-MDP starts after the transition to new fixation $\mathbf{x}_{k+1}$.

The information gain that would be obtained, at step $k+1$, by candidate fixation $\overline{\mathbf{x}}$ is

$$\Delta I(\mathbf{b}_k, \overline{\mathbf{x}}) = -\left(H(\mathbf{b}_{k+1}) - H(\mathbf{b}_k)\right) \tag{9}$$

$$\text{where} \quad H(\mathbf{b}_k) = -\sum_{\ell=1}^{L} b_{k,\ell}\, \log(b_{k,\ell}) \tag{10}$$

is the Shannon entropy. The entropy represents the observer's uncertainty about the target location. The Infomax policy chooses the fixation $\overline{\mathbf{x}}$ that maximizes the expected information gain

$$\pi: \quad \overline{\mathbf{x}}_{k+1} \leftarrow \arg\max_{\overline{\mathbf{x}}} E[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})]. \tag{11}$$

Hence the choice of location, according to the Infomax model, minimizes the expected Shannon entropy of the belief state, i.e. uncertainty of target location. The expected information gain $E[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})]$ is calculated analytically by Najemnik and Geisler[11] as

$$E[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})] = \frac{1}{2} \sum_{\ell=1}^{L} b_{k,\ell}\, V^2(\mathbf{u}_\ell - \overline{\mathbf{x}}). \tag{12}$$

Figure 1 illustrates the decision-making process of the Infomax model.

2.2. *Parametrization of visibility map*

The visibility maps $V(\boldsymbol{x})$ are represented by parametric models $V_{\boldsymbol{\theta}}(\mathbf{x})$. These comprise sums of polar basis functions, with coefficients $\boldsymbol{\theta}$, as follows:

$$V_{\boldsymbol{\theta}}(\mathbf{x}) = V_{\boldsymbol{\theta}}(r, \psi) =$$

$$\left| \sum_{m=0}^{M} \sum_{n=0}^{N} \left(\vartheta_{m,n} B_{m,n}(r, \psi) + \vartheta_{m,n}^\star B_{m,n}^\star(r, \psi)\right) \right| + V_\epsilon$$

$$\text{where} \quad r = \|\mathbf{x}\|, \quad \psi = \tan^{-1}(y, x),$$

$$\boldsymbol{\theta} = \left(\vartheta_{0,0}, \ldots, \vartheta_{M,N}\,; \vartheta_{0,0}^\star, \ldots, \vartheta_{M,N}^\star\right). \tag{13}$$
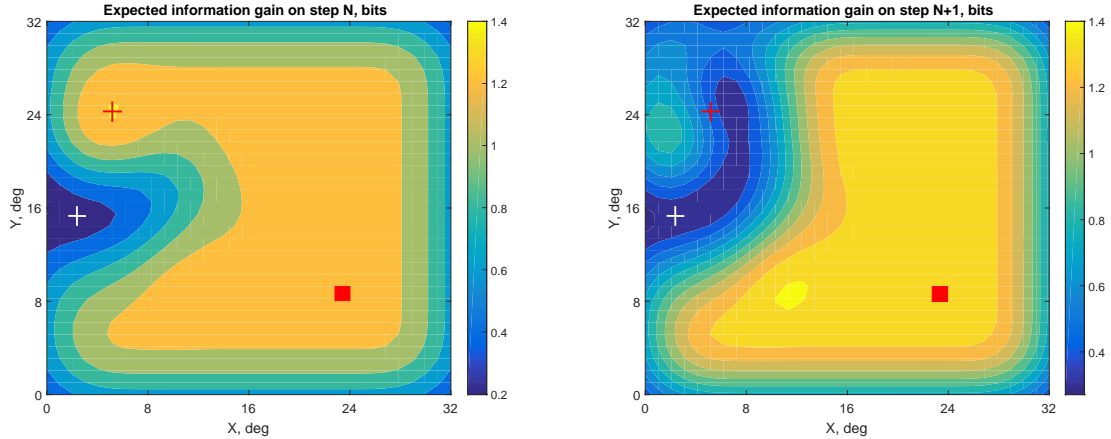
Fig. 1.   The decision making process in Infomax model. The left map represents the expected information gain (eqn. 12), across the visual field. The white cross corresponds to the location of the current fixation on step k, while the red square corresponds to the location of the target. The observer makes a decision to fixate at the location defined by policy: $\overline{\mathbf{x}}_{k+1} \leftarrow \pi(\mathbf{b}_k)$. This decision results in the saccadic eye-movement to location $\mathbf{x}_{k+1} = \alpha(\overline{\mathbf{x}}_{k+1})$ marked by the red cross. After receiving the observation at step $k+1$, the observer updates its belief state. In this particular situation, the target is not in the vicinity of $\mathbf{x}_{k+1}$, and the action resulted in reduced probability b_{k+1} around the fixation (red cross).

Here $B_{m,n}(r, \psi)$, and $B^*_{m,n}(r, \psi)$ are cosine and sine Bessel radial functions, respectively (see appendix A.1 for details). The parameter vector $\boldsymbol{\theta}$ contains the corresponding $2MN$ coefficients, $\vartheta_{m,n}$ and $\vartheta^*_{m,n}$. The small regularization constant $V_\epsilon = 0.01$ is used to avoid any singularity in the precision parameter of the observation model (3). Recall that the visibility map $V(\mathbf{x})$, which represents the signal-to-noise ratio across the visual field, is a positive function. For this reason, the absolute value of the sum is taken in (13). Figure 2 shows some example visibility maps, with parameters $\vartheta_{0,1} = 30$, $\vartheta_{1,1} = -30$, $R_0 = 8°$.

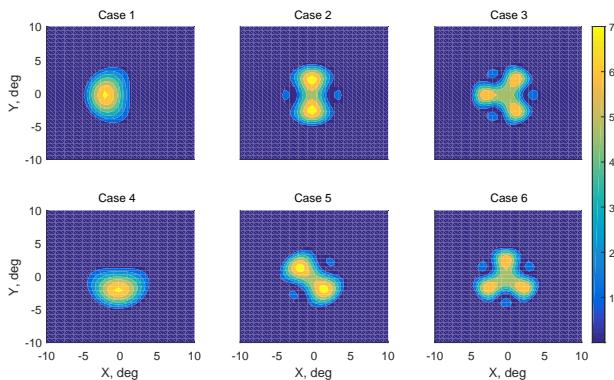

Fig. 2.   Examples of visibility map with parameters: $\vartheta_{1,0} = 30$ (shared for all examples), $\vartheta_{1,1} = -30$; $\vartheta_{1,2} = -30$; $\vartheta_{1,3} = -30$ (upper row); $\vartheta^*_{1,1} = -30$; $\vartheta^*_{1,2} = -30$; $\vartheta^*_{1,3} = -30$ (lower row) and zeros for all other parameters.

2.3. *Asymmetry of saccade distribution*

In this section, we show how the asymmetry of the visibility map function influences the distribution of eye-movement vectors in the Infomax model. We simulated the Infomax model on a computational grid of size 48×48, and varied the parameters $\boldsymbol{\theta}$ of the visibility map function. We fixed the number of radial and angular dimensions to $M = 1$ and $N = 5$, while the radius of the visual field is $R_0 = 8°$. The VMs were generated using a special case $\boldsymbol{\theta}(\alpha, \beta, \ell)$ of the parameter vector (13), with entries

$$\vartheta_{m,n} = \begin{cases} \alpha & \text{if } n = 0 \\ \alpha\beta & \text{if } n = \ell \\ 0 & \text{otherwise} \end{cases} \qquad \vartheta^\star_{m,n} = 0. \qquad (14)$$

The parameters α and β regulate the overall level of visibility and angular amplitude respectively, for a given frequency ℓ. Parameters were selected as: $\alpha \in \{30, 60\}$, $\beta \in \{0, 0.1, ..., 1\}$, $\ell \in \{1, 2, 3, 4, 5\}$. For each visibility map, defined by the above parameters, we run 10000 episodes of the PO-MDP model, and calculate the histogram of saccade vectors. For each case we measured the number of saccades in the directions that correspond to regions of suppressed and enhanced of visibility (see Fig. 3, for example). We define the tendency to move the eyes in the scotoma direction as a ratio: the number of saccades towards the suppressed region, divided by the number towards the enhanced region.
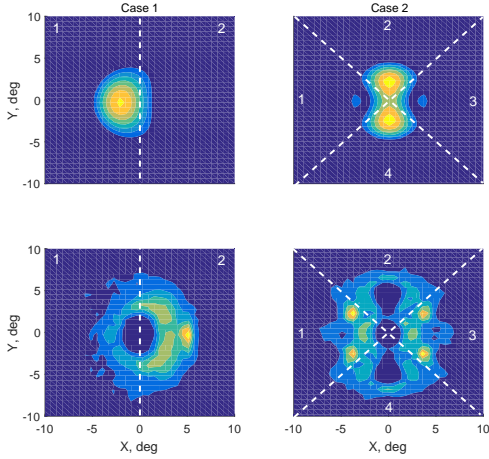
Fig. 3. The visibility maps (top) and histograms of saccade vectors (bottom) for parameter cases: $\alpha = 60$, $\beta = 1$, $\ell \in \{1, 2\}$. For the case $\ell = 1$ (left), the left hemifield is the direction of enhanced visibility, marked by number 1. We observe the deficit of saccade towards this direction; on the opposite side there is a surplus, in direction 2. For the case $\ell = 2$ (right) we found a surplus of saccades towards directions 1 and 3, which correspond to suppressed visibility, and a deficit towards directions 2 and 4.

We find that in the low visibility case of $\alpha = 30$, the cases $\ell \in \{1, 3\}$ show monotonic growth of preference towards the scotoma direction, as a function of a parameter β (Fig. 4). In the cases $\ell \in \{2, 4, 5\}$, we observed that the direction of better visibility is preferred for most values of parameter β. For the high visibility case $\alpha = 60$, with $\ell \in \{1, 2, 3, 5\}$, we found that the deficit direction is more preferred. This preference grows monotonically with β, for $\ell \in \{1, 2, 3\}$, as shown in Fig. 5. In the case $\ell = 4$ the enhanced direction is less preferred, but the difference is less than 5% of the total saccade count.
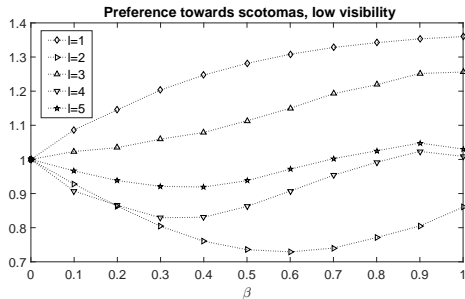

Fig. 4. Preference towards scotoma direction in the low-visibility case $\alpha = 30$. The cases $\ell \in \{1, 3\}$ show monotonic growth of preference towards the scotoma direction with parameter β.

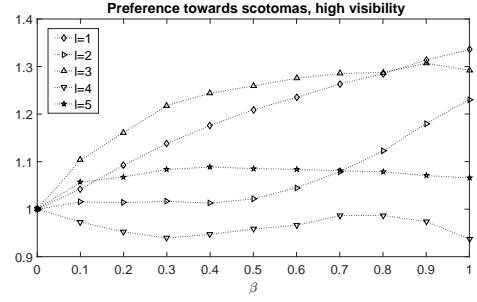

Fig. 5. Preference towards scotoma direction in the high-visibility case $\alpha = 60$. In cases $\ell \in \{1, 2, 3, 5\}$ the deficit direction is preferred, and this preference grow monotonically with β for cases $\ell \in \{1, 2, 3\}$.

On the basis of our findings, we conclude that a preference towards scotoma direction arises from the Infomax model, given a high general level, and low angular frequency of the visibility map. In the case of low visibility, we find the preference towards the scotoma direction only for the first two odd angular terms. Due to the weakening of bias towards the deficit direction with angular number ℓ, we limit the number of angular dimensions to $N = 3$, in our simulations.

Interestingly, the proportion of saccades towards scotoma direction for the case $\ell = 1$, $\beta = 1$, which corresponds to the situation of complete blindness in deficit hemifield, is consistent with that reported for patients with homonymous hemianopia (visual field loss on the same side, for both eyes) performing a visual search task.[32]

3. Parameter learning algorithm

This section develops an algorithm for estimating the parameters of the underlying visibility map, given the observed eye-movement sequences and the experimental task.

3.1. *Probability of fixation sequences*

An individual scan-path, corresponding to one experimental trial, comprises a sequence of fixations $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_k, \ldots, \mathbf{x}_K)$. We define a *data set* as the concatenation $\mathcal{X}$ of S observed scan-paths:

$$\mathcal{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_S). \qquad (15)$$

Recall that we have defined the visibility map through parameter set $\boldsymbol{\theta}$, in sec. A.1. Our goal is to maximize the log-probability of a given data set $\mathcal{X}$, given a model

defined by parameters $\boldsymbol{\theta}$:

$$\boldsymbol{\theta}^{\star} = \arg\max_{\boldsymbol{\theta}} \; \log \, p(\mathcal{X}|\boldsymbol{\theta}). \tag{16}$$

We use an optimization method that is inspired by the *Policy Gradient with Parameter Exploration* scheme of Sehnke et al.[16,17] This requires a probability distribution $p(\boldsymbol{\theta}|\boldsymbol{\rho})$ to be defined over possible models $\boldsymbol{\theta}$. Then we can optimize the following objective function, with respect to the upper-level parameters $\boldsymbol{\rho}$:

$$\log p(\mathcal{X}|\boldsymbol{\rho}) = \int_{\boldsymbol{\Theta}} \log\big(p(\mathcal{X}|\boldsymbol{\theta})\big) \, p(\boldsymbol{\theta}|\boldsymbol{\rho}) \, \mathrm{d}\boldsymbol{\theta}. \tag{17}$$

The gradient of this log-probability, with respect to the upper-level parameters $\boldsymbol{\rho}$ can be expressed as follows:

$$\nabla_{\boldsymbol{\rho}} \log\big(p(\mathcal{X}|\boldsymbol{\rho})\big) = $$
$$\int_{\boldsymbol{\Theta}} \log\big(p(\mathcal{X}|\boldsymbol{\theta})\big) \, p(\boldsymbol{\theta}|\boldsymbol{\rho}) \, \nabla_{\boldsymbol{\rho}} \log p(\boldsymbol{\theta}|\boldsymbol{\rho}) \, \mathrm{d}\boldsymbol{\theta} \tag{18}$$

where the 'log likelihood trick' $\nabla_{\boldsymbol{\rho}} \log p(\boldsymbol{\theta}|\boldsymbol{\rho}) \equiv \nabla_{\boldsymbol{\rho}} p(\boldsymbol{\theta}|\boldsymbol{\rho})/p(\boldsymbol{\theta}|\boldsymbol{\rho})$ is used to express the derivative of the factor $p(\boldsymbol{\theta}|\boldsymbol{\rho})$ in (17). The gradient integral (18) can now be approximated, by sampling J parameter vectors $\boldsymbol{\theta}_j$ from the distribution $p(\boldsymbol{\theta}|\boldsymbol{\rho})$, and computing the sum

$$\nabla_{\boldsymbol{\rho}} \log\big(p(\mathcal{X}|\boldsymbol{\rho})\big) \approx \frac{1}{J} \sum_{j=1}^{J} \log\big(p(\mathcal{X}|\boldsymbol{\theta}_j)\big) \nabla_{\boldsymbol{\rho}} \log p(\boldsymbol{\theta}_j|\boldsymbol{\rho})$$
$$\text{where} \quad \boldsymbol{\theta}_j \sim p(\boldsymbol{\theta}|\boldsymbol{\rho}). \tag{19}$$

For each parameter set $\boldsymbol{\theta}_j$, we evaluate the probability $p(\mathcal{X}|\boldsymbol{\theta}_j)$ according to section A.3.

3.2. *Estimation algorithm*

The set of upper-level parameters $\boldsymbol{\rho}$, which is used to define the parameter distribution $p(\boldsymbol{\theta}|\boldsymbol{\rho})$, comprises a mean and standard deviation for each of the $2MN$ visibility parameters in (13), as follows:

$$\boldsymbol{\mu} = \big(\mu_1, \ldots, \mu_{2MN}\big), \quad \boldsymbol{\sigma} = \big(\sigma_1, \ldots, \sigma_{2MN}\big)$$
$$\text{and} \quad \boldsymbol{\rho} = (\boldsymbol{\mu}, \boldsymbol{\sigma}). \tag{20}$$

Each parameter θ_i is sampled J times from the corresponding normal distribution $p(\theta_i|\boldsymbol{\rho})$, such that $\boldsymbol{\theta}_j$ is a complete sample vector, with entries

$$\theta_{i,j} \sim N(\mu_i, \sigma_i), \quad i = 1, \ldots, 2MN. \tag{21}$$

The learning of the upper-level parameters $\boldsymbol{\rho}$ is an iterative process, which involves estimating the gradient of log-probability (17), via (19), and then updating the upper-level parameters $\boldsymbol{\rho}$; this constitutes one *learning*

epoch. We define the convergence criterion for the gradient as: $\|\nabla_{\boldsymbol{\rho}} \log(p(\mathcal{X}|\boldsymbol{\rho}))\| < \varepsilon$, where ε is chosen as $1/25$-th of the initial gradient magnitude.

In practice, a *symmetric sampling* scheme is used for (21), because it provides more robust gradient approximation.[16] The complete algorithm is presented in Fig. 6.

INITIALIZATION

    (1) Construct a data set $\mathcal{X}$, by concatenating multiple scan-paths, as in (15).

    (2) Initialize the upper-level parameters $\boldsymbol{\rho}$, as $\sigma_i \leftarrow 1$, and $\mu_i \sim N(0,1)$, for all $i = 1, \ldots, 2MN$.

LEARNING

while $\|\nabla_{\boldsymbol{\rho}} \log(p(\mathcal{X}|\boldsymbol{\rho}))\| > \varepsilon$

    (1) Sample $2J$ parameter sets $\boldsymbol{\theta}_j^{\pm} = \boldsymbol{\mu} \pm \boldsymbol{\epsilon}_j$, as in sec. A.2, where $\boldsymbol{\epsilon}_j \sim N\big(\mathbf{0}, \mathrm{diag}(\boldsymbol{\sigma})\big)$.

    (2) Estimate probabilities $p(\mathcal{X}|\boldsymbol{\theta}_j^{\pm})$ of data set $\mathcal{X}$, as in sec. A.3.

    (3) Compute gradient $\nabla_{\boldsymbol{\rho}} \log(p(\mathcal{X}|\boldsymbol{\rho}))$ according to eqns. (A.4,A.5).

    (4) Update upper-level parameters $\boldsymbol{\rho}$ according to:

$$\boldsymbol{\rho} \leftarrow \boldsymbol{\rho} + \lambda \nabla_{\boldsymbol{\rho}} \log(p(\mathcal{X}|\boldsymbol{\rho}))$$

    Where λ is the learning rate constant.

OUTPUT

$\boldsymbol{\theta} \leftarrow \boldsymbol{\mu}$

Fig. 6.   The visibility map inference algorithm, using the symmetric sampling scheme of Sehnke et al.[16]

3.3. *Performance on simulated data*

Here we demonstrate the performance of the proposed inference algorithm, on eye-movement trajectories simulated by the Infomax model. The set of ground truth visibility maps was generated by random selection of each observation parameter θ_i in the interval $[-1, 1]$. For each visibility map function in this set we ran 100 episodes of PO-MDP on a 64×64 computational grid. These generated trajectories served as input for the inference algorithm. The inferred visibility maps were then compared with ground truth visibility maps. Each input sequence is split between the training and vali-

dation sets in a ratio of 75/25. The training set is used for evaluation of observation parameters. The validation set is used for assessment of the performance of learning algorithm.

We set the learning rate as $\lambda = 0.5$ in all cases. On each computational epoch we compute the following performance characteristics:

- Euclidean distance between ground truth VM and inferred VM (Figure 7) on the current computational epoch divided by Euclidean norm of ground truth.

$$R = \frac{\sqrt{\sum_{\ell=1}^{L}\big(V_{\text{est}}(\mathbf{u}_\ell) - V_{\text{true}}(\mathbf{u}_\ell)\big)^2}}{\sqrt{\sum_{\ell=1}^{L} V_{\text{true}}^2(\mathbf{u}_\ell)}} \qquad (22)$$

- The log-probability of single saccade for current parameter values for training (Figure 8) and validation sets (Figure 9) compared to chance level:

$$\log(p_s) = \frac{\log\big(p(\mathcal{X}, \boldsymbol{\theta})/D\big)}{\text{length}(\mathcal{X})} \qquad (23)$$

where $D = 64 \times 64$ is the dimensionality of the grid. This representation is used because the length of eye-movement sequence is different for each patient.

- Jensen-Shannon distance between the spatial distributions of input and simulated action sequences (Fig. 10).

For these parameter settings, the algorithm reaches the convergence condition at the 25th computational epoch, according to both Euclidean and Jensen-Shannon distances. We conclude that the algorithm simultaneously satisfies both of these performance criteria. However, the ground truth visibility map cannot be inferred exactly, due to the variance of the estimate A.3. Figure 11 shows three examples of inferred visibility maps (lower row) from simulated eye-movements with target visibility maps (upper row).
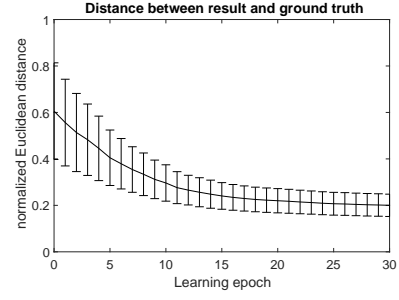

Fig. 7. The Euclidean distance between ground truth and estimated visibility map, divided by the Euclidean norm of the ground truth, for current values of the parameters on each computational epoch.
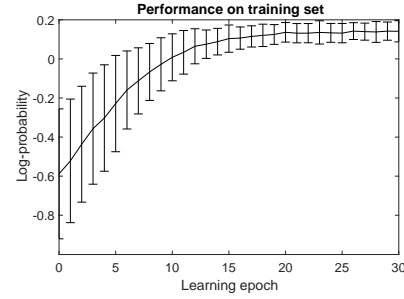

Fig. 8. The log-probability of a single saccade, in the training set, for each computational epoch.
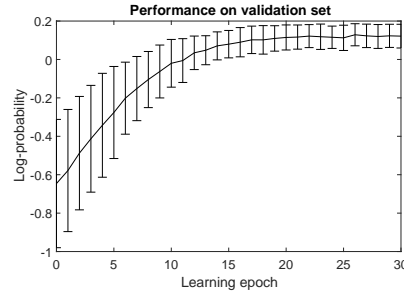

Fig. 9. The log-probability of a single saccade, in the validation set, for each computational epoch.
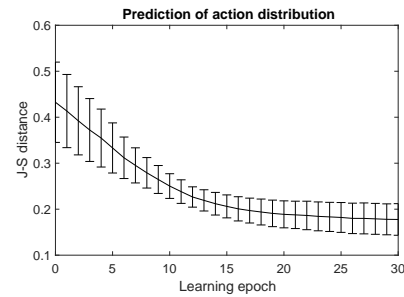

Fig. 10. Jensen Shannon distance between the spatial distribution of input and simulated action sequences.
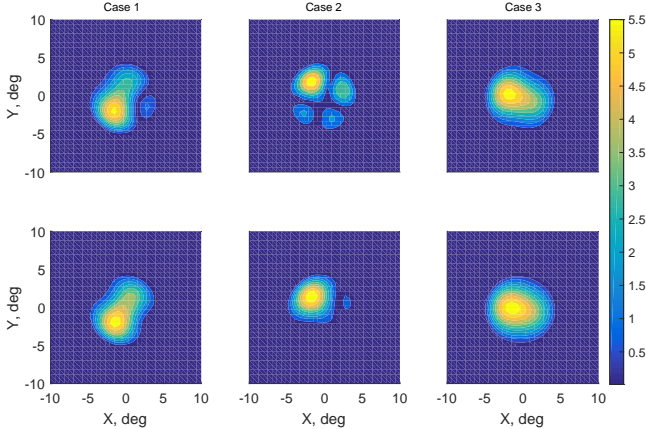
Fig. 11.   The inferred visibility maps (lower row) from simulated eye-movements, with the target visibility maps (upper row) that correspond to three parameter cases.

In the appendix A.4 we examine the influence of the length of the input trajectory on the performance of the inference, and confirm that the lengths of both simulated and experimental trajectories are sufficient for estimation of target VM with precision that is close to the optimal one.

4.   Analysis of clinical AMD data

In this section we compare our estimates of the visibility maps with the values of visual field tests provided to us by Janssen et al.[3] and Van der Stigchel et al.[2]

4.1.   *Summary of datasets*

We have used the following two sets of eye-tracking data, to evaluate the inference algorithm.

Dataset 1. Eye tracking data from a psychophysical experiment of Van der Stigchel et al.[2] In this experiment four AMD patients performed a visual search task. The stimulus consisted of a visual search field with 32 C-shaped distractors, and a single O-shaped target. The target and distractors were positioned on a hexagonal lattice. The positions of the target and distractors, and the orientation of the distractors, were randomized on each trial. Each participant completed 55 trials of visual search. An Eyelink 1000 infrared eye-tracker was used to record of eye-movements, and visual field values are provided for all four AMD patients.

Dataset 2. Eye tracking data from a visual search experiment of Janssen et al.[3] In this experiment, 12 AMD patients completed 10 to 15 trials per session, in which they searched for target objects in a scene. The target consisted of a two-dimensional Gaussian blob, with a spatial deviation of 0.5 deg. The stimulus backgrounds were randomly selected from a database of 80 images of outdoor scenes. There were between zero and nine target objects, which were superimposed at random locations in the given image. The number of targets was chosen randomly in each trial. Visual field values were provided for all 12 AMD patients.

If the landing position of the last saccade is within $2°$ of the target, then the trial is considered successful, for eye-movement trajectories from all two data-sets. Unsuccessful trials were excluded from consideration.

4.2.   *Performance on experimental data*

In this subsection we demonstrate the performance of the inference algorithm on the datasets from psychophysical experiments.[2,3]

The eye-movement data comprises sequences $\mathbf{X}_s$ of fixations, which are binned into a computational grid of resolution 64×64. This representation of data is used during the training phase and the evaluation of performance. For each patient, we split the entire eye-movement time-series into the training and test sets, in the ratio of 75/25. The training set was used to evaluate the observation parameters. The validation set was used for assessment of the performance of the learning algorithm. On each learning epoch we evaluated the following metrics: average Jensen-Shannon distance between target distribution and simulated eye-movements and average log-probability of single saccade, compared to the chance level (23).
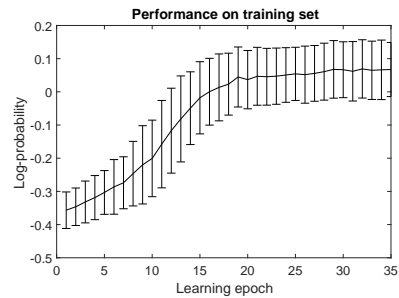

Fig. 12.   The log-probability of a single saccade in the training set for each computational epoch.
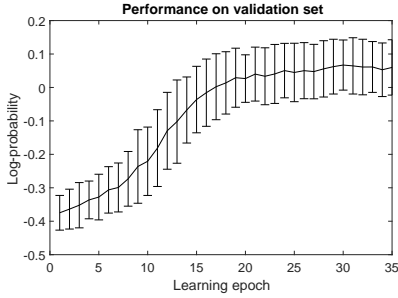
Fig. 13.   The log-probability of a single saccade in the validation set for each computational epoch.
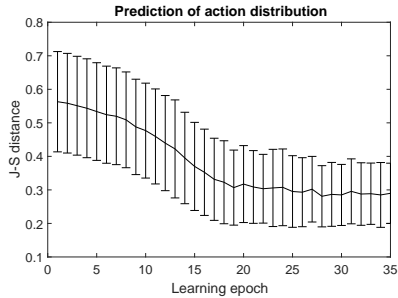

Fig. 14.   Jensen-Shannon distance between the distributions of simulated and experimental eye-movements

We set the learning rate as: $\lambda = 0.5$ in this test, and in the evaluation of VMs in section 4.4. The log-probability of a single saccade, $\log(p_s)$, was averaged over all AMD patients and initial conditions. The figures 12 and 13 show the average log-probability of a single saccade as a function of learning epoch in the training and validation sets correspondingly. In most cases, the log-probability value stabilizes around the 30th learning epoch. We show the Jensen-Shannon distance[33] between the experimental and simulated saccade distributions in Fig. 14.

### 4.3. *Results of inference for dataset 1*

In this subsection we discuss the results of the inference algorithm applied on eye-tracking data from Van der Stigchel et al.[2] Firstly, we present the histograms of saccade vectors in figure 16. The upper and lower rows of the figure plot the experimental and the simulated saccade distributions, respectively. Similar to the experiment of Janseen et al., the distributions of the saccade vectors are anisotropic, and have maxima at the locations away from the fovea. The main difference in the patients' behaviour in this experiment, compared to the previous one, is a stronger bias towards the horizontal direction. The average Jensen-Shannon distance

between experimental and simulated distributions of the saccade vectors was 0.24, compared to the value of 0.22 for the Janssen et al. data.

The figure 18 presents the inferred VMs and the visual field values for AMD patients. We can see that the algorithm has identified lower sensitivity locations which are consistent with visual field values. In the case of P1, the algorithm identified the asymmetry of the visibility map and shifted the location of maximal sensitivity to the left hemifield. In the case of P2, the impaired fovea and right hemifield were identified by our inference scheme. In the case of P4, we estimated a distribution of sensitivity that is similar to the case of P2. The case of P3 corresponds to an irregular distribution of the scotomas. The algorithm estimated the symmetric VM with low values of sensitivity in the foveal region.

4.4. *Results of inference for dataset 2*

Figure 15 shows the distribution of simulated fixations at the end of the learning process, for simulated agents.[3] The intersection of dotted lines marks the location of maximum probability. Note that, in general, the probability density function is asymmetric, and has a maximum away from the fovea. The Infomax model, using the learned visibility map, captures this property for almost all observers. The exception is patient P5, who had a significantly shorter data-set of eye-movement trajectories, due to better visual conditions[3] and, therefore, higher performance in task execution. The table 4.4 shows three performance metrics of inference algorithm: the distance between locations of maxima of experimental and simulated saccade distributions, Jensen-Shannon distance between predicted and experimental distributions, and the p-values calculated according to χ^2 test statistics. The patients are numbered as P2–P11 for the Janssen et al. experiment[3] and P1d–P4d for the van der Stigchel et al. experiment.[2] For most patients the Jensen-Shannon distance between model and experimental distributions is lower than 0.3, and the location of a maximum was predicted with error lower than 3°. Figure 17 shows the binocular scotoma maps[3] (left) and inferred values of VMs (right). The binocular scotoma maps are represented as detection rate of the target object across the visual field. The scotoma maps were measured for two sizes of object: $5° \times 5°$ and $2.5° \times 2.5°$. For patient P10 the measurements were provided for a stimulus size of $6.5° \times 6.5°$. One can see in Figure 17 (right) that the
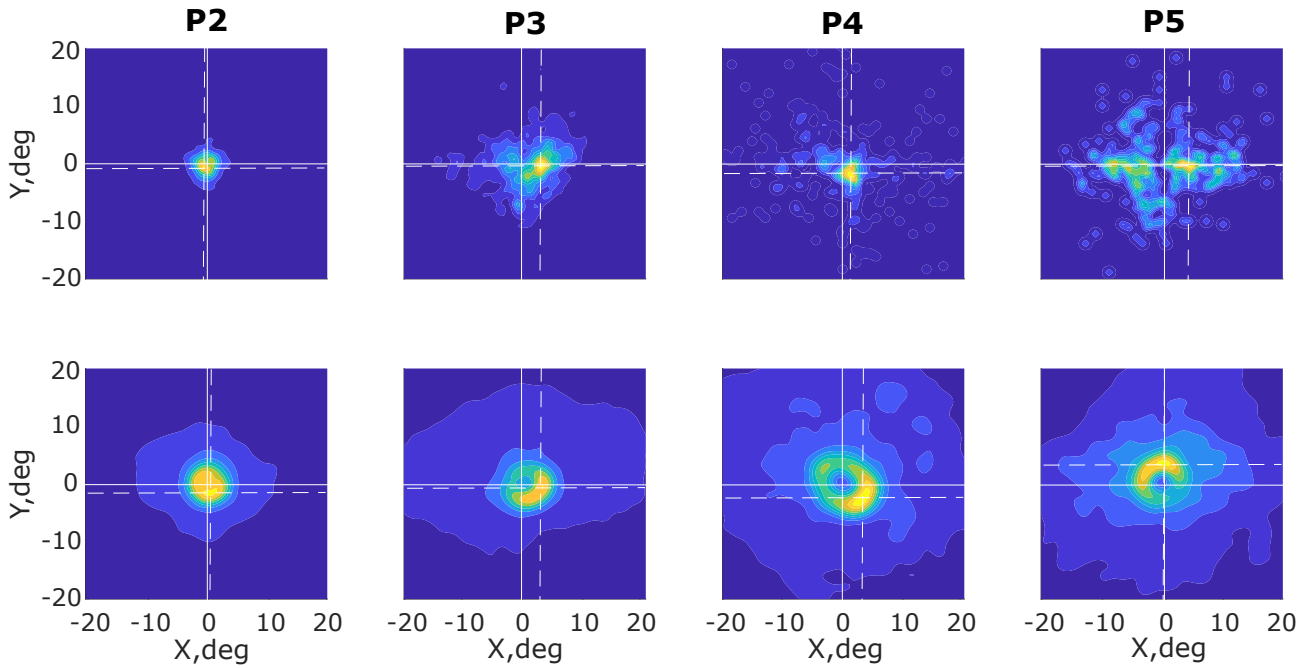
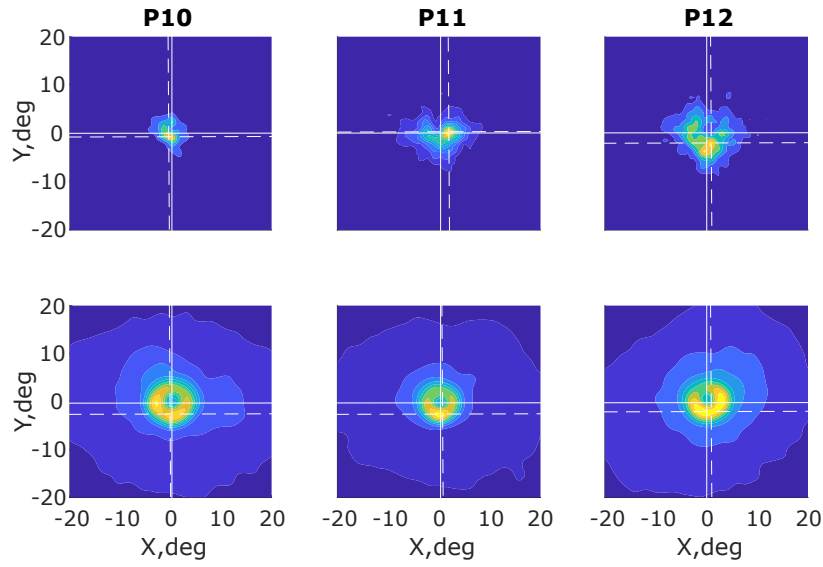(a) Results for patients P2–5 in the the Janssen et al. data set.[3]
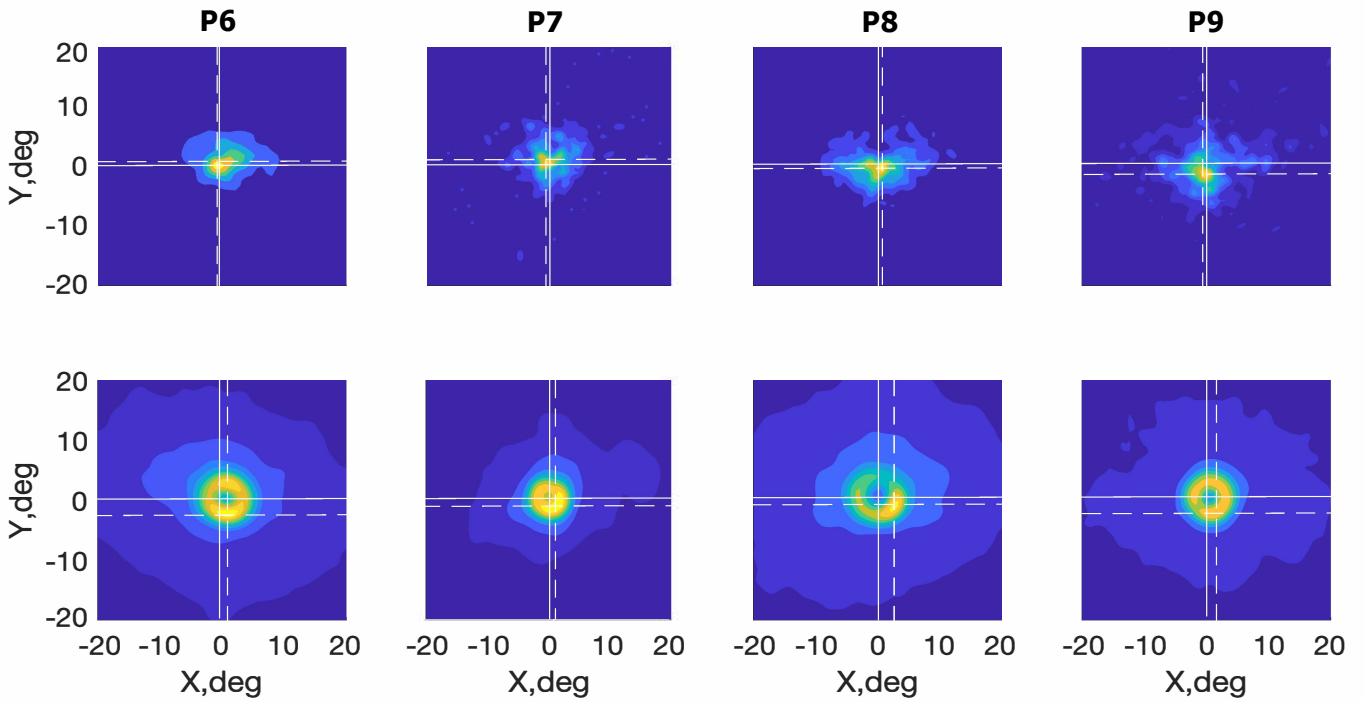
(b) Results for patients P6–9 in the Janssen et al. data set.[3]

Fig. 15.   The distribution of experimental eye-movements (upper row) and distributions of simulated actions (lower row), with axes are in degrees of visual angle. The intersection of dotted lines marks the maximum of each probability density function.

inferred visibility maps are irregular, unlike those of the normal controls performing a visual search task.[11]

We assume that the differences between the VMs are caused by the presence of the scotomas. The target de-

(a) Results for patients P10–12 in the Janssen et al. data set.[3]

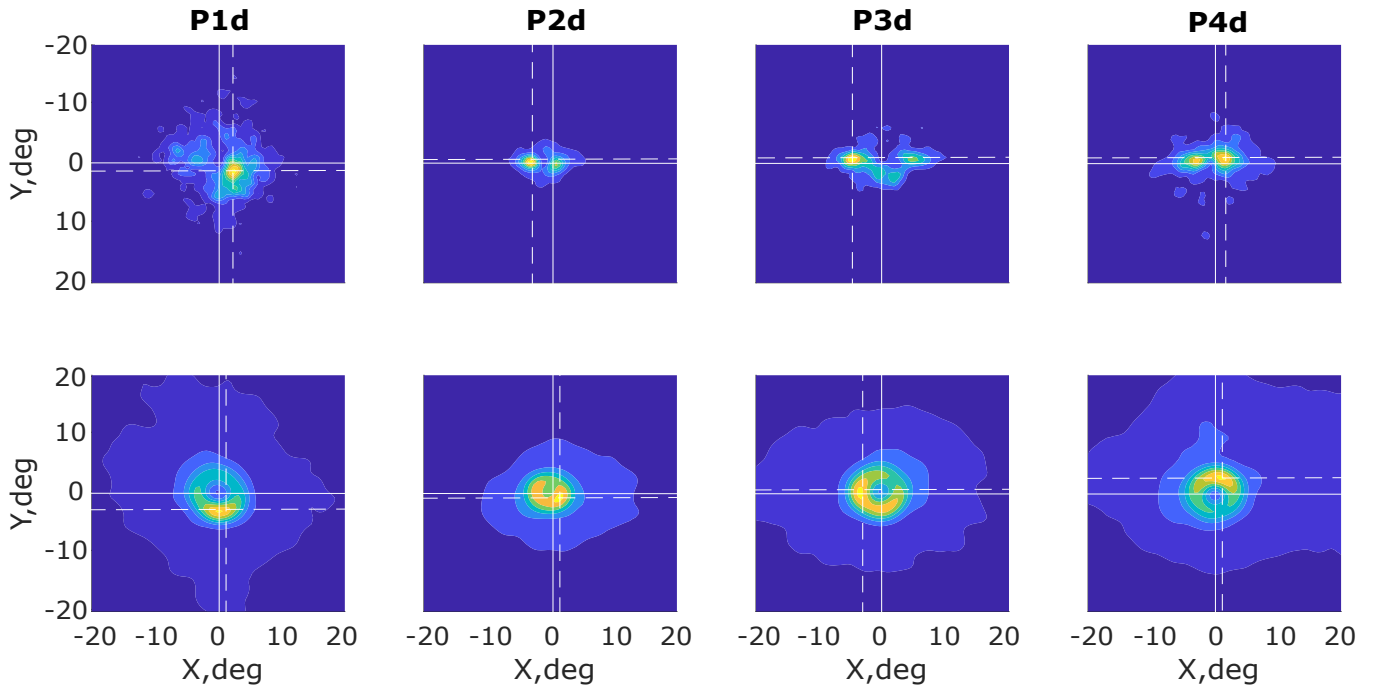(b) Results for patients P1–4 in the van der Stigchel et al. data set.[2]

Fig. 16. The distribution of experimental eye-movements (upper row) and distributions of simulated actions (lower row), with axes in degrees of visual angle. The intersection of dotted lines marks the maximum of each probability density function.

tection rate is close to zero in the dark blue areas of the images in Fig. 17 (right). We associate these locations with the periphery if the distance between them and fovea is larger than 5°, and otherwise with the scotomas. For the patients P4, P5, P6, P7, P11 and P12, the inference algorithm identified the scotomas in the quadrants consistent with binocular scotoma maps 17. However, even for these cases with asymmetrical visual field values, we can't clearly see the asymmetrical inferred VMs. This happens for several reasons.
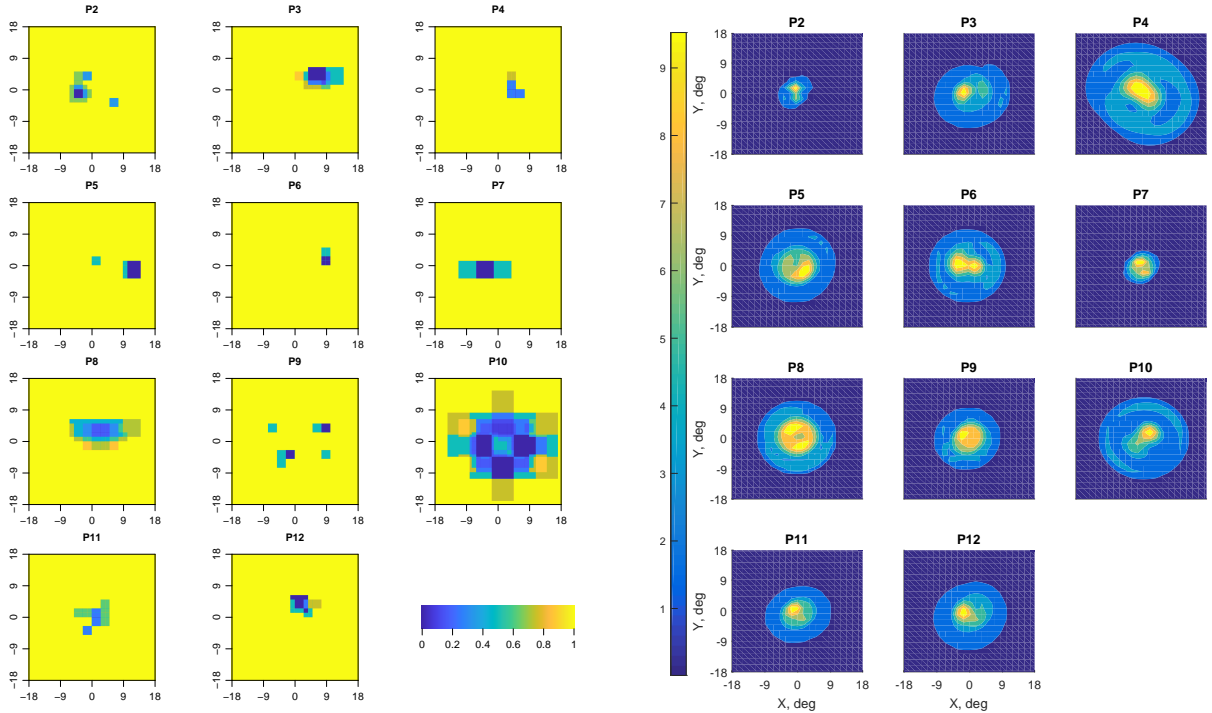
(a) The binocular scotoma maps for AMD patients in the Janssen et al. experiment.[3] The scotoma maps were measured for three sizes of target: $5° \times 5°$, $2.5° \times 2.5°$ and $6.5° \times 6.5°$. For patient P10 the measurements were provided for stimulus size of $6.5° \times 6.5°$.

(b) The inferred values of visibility maps for AMD patients performing the visual search task in the Janssen et al. experiment.[3] See Fig. 18 (right) for the visibility maps from the van der Stigchel et al. experiment.[2]

Fig. 17.

First of all, our algorithm has limited precision even in the case of simulated eye-movements 3.3. Secondly, we limited the number of angular dimensions to $N = 3$ as described in section 2.3. The limited dimensionality is also a reason why the distributions of simulated eye-movements 16 do not reflect the irregularity of the experimental distributions. The third reason is a difference in conditions during the measurement of visual field values and VMs. In the search experiments, VMs are measured in the presence of background noise, which results in lower performance in the periphery. This is not true for visual field values that do not decline with eccentricity.
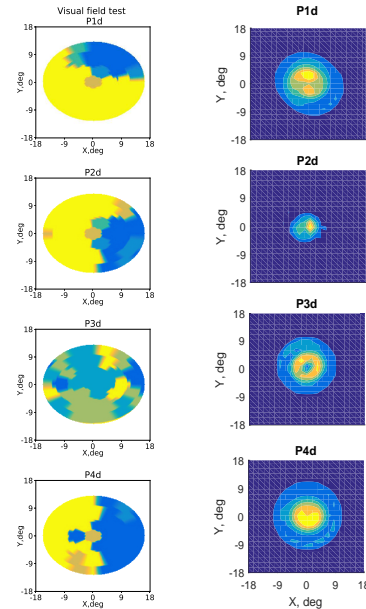

Fig. 18.   The results of the visual field test for AMD patients (left) and the inferred visibility maps (right) for the van der Stigchel et al. experiment.[2]

Case	Localization error, deg.	J-S distance	p-value
P2	1.67	0.214	0.875
P3	2.00	0.178	0.764
P4	1.00	0.235	0.863
P5	2.33	0.345	0.692
P6	2.00	0.256	0.845
P7	1.33	0.397	0.623
P8	1.33	0.145	0.835
P9	1.66	0.169	0.824
P10	1.33	0.226	0.792
P11	1.00	0.239	0.824
P12	2.00	0.163	0.889
P1d	1.33	0.228	0.778
P2d	1.00	0.246	0.882
P3d	3.00	0.237	0.856
P4d	2.33	0.233	0.874

Table 4.4 Prediction accuracy for the spatial distributions. The patients from data set[2] are marked by letter 'd' to distinguish from the previous patients.

4.5. *Identification of scotoma regions*

Our main objective was to determine how the presence of the scotomas influences visual search behaviour for AMD patients. In this section we study the correlation of four variables: visual field values, inferred visibility maps, experimental distributions and simulated distributions. The values of these four variables were integrated in each of the four quadrants. The resulting vectors of the experimental and the simulated distributions of saccades were normalized for each patient. These vectors are referred to as *quadrant distributions*.

We computed the Jensen-Shannon distances between variables for each patient and averaged the results over patients (see table 4.5 below diagonal). This measure is called intra-patient statistical distance. In order to evaluate the performance of the algorithm we estimated the Jensen-Shannon distance between variables across different patients (inter-patient statistical distance). For each variable and patient we made 14 comparisons (15 begin the number of patients), and averaged the results over all comparisons $\frac{1}{N}\sum_{j=1}^{N} \mathrm{JS}_j$. The result was averaged again over all patients giving the average Jensen-Shannon distance for each variable (see 4.5 above diagonal).

In order to calculate the variance of estimates of average statistical distance we used the jackknife method of resampling.[34] We estimated the average JS distance omitting one comparison:

$$\overline{\mathrm{JS}}_i = \frac{1}{N-1}\sum_{j=1,j\neq i}^{N} \mathrm{JS}_j \qquad (24)$$

On the next step we estimate the standard deviation of the jackknife estimator:

$$\sigma(\overline{\mathrm{JS}}) = \sqrt{\frac{N-1}{N}\sum_{i=1}^{N}\left(\overline{\mathrm{JS}}_i - \overline{\mathrm{JS}}\right)^2} \qquad (25)$$

where $\overline{\mathrm{JS}} = \frac{1}{N}\sum_i^N \overline{\mathrm{JS}}_i$. Next, we average the standard deviation of the jackknife estimator over all patients. The results are reported in the table 4.5.

Variable	VF values	Inferred VM	Exp. distr.	Sim. distr.
Visual field values	—	0.25 ± 0.04	0.46 ± 0.05	0.49 ± 0.06
Inferred VM	0.19	—	0.43 ± 0.06	0.39 ± 0.05
Experimental distr.	0.43	0.43	—	0.15 ± 0.05
Simulated distr.	0.47	0.38	0.09	—

Table 4.5 Average JS distance between variables for each patient below the diagonal, and average JS distance between variables across different patients with standard errors above the diagonal

We can see that the difference between the intra-patient and the inter-patient statistical distance is significant only for two cases: the distance between inferred VM and VF values, and the distance between the simulated and the experimental distributions of eye-movements. For all other cases there is a chance level of similarity both for the intra-patient and the inter-patient comparisons. The inter-patient statistical distance is larger than the intra-patient distance for all cases.

5. Conclusion

We have developed and tested an algorithm for the inference of visibility maps of both human observers and computational agents.

On the basis of the average Jensen-Shannon distance at the end of the learning process, the prediction error of maximum location, and the visual similarity between experimental and simulated distributions, we conclude that our Infomax model is capable of describing the spatial distribution of eye-movements for AMD

patients. Interestingly, we find that neither the Infomax model nor the human observer display a preference towards the scotoma direction, which was observed only in two experimental cases.[2] We observe that a preference towards impaired regions of the visual field occurs only in the most trivial configurations of VM, such as in the case of homonymous hemianopia (visual field loss on the same side, for both eyes). In this case, the ratio of the saccades towards impaired regions can be calculated exactly, using the Infomax simulations.

Our algorithm is an iterative procedure, which evaluates the gradient of fixation log-probability, and subsequently updates the observation parameters. The algorithm was inspired by the reinforcement learning method of parameter exploring policy gradient.[16] Despite the simplicity of this gradient-based approach, it shows good performance, compared to alternative methods, in the reconstruction of the observation model for both computational agents and human observers. Our algorithm can easily be generalized for any Po-Mdp with independent observation parameters and a convolutional policy.[23] In the future we plan to extend this approach for generative models of eye-movements that are based on reinforcement learning methods.[8,13,23]

Acknowledgements

Appendix A
Implementation of learning algorithm

This section gives full details of the parameter estimation algorithm, including the VM parameterization, the sampling scheme, and the probability computations.

A.1. *Bessel Circular Functions*

We use Bessel Circular Functions[35] for parametrization of the visibility map:

$$B_{m,n}(r,\psi) = \begin{cases} N_{mn} J_n(c_{m,n}r/R)\cos(n\psi) & \text{if } r < R \\ 0 & \text{otherwise} \end{cases}$$
(A.1)

$$B^*_{m,n}(r,\psi) = \begin{cases} N_{mn} J_n(c_{m,n}r/R)\sin(n\psi) & \text{if } r < R \\ 0 & \text{otherwise} \end{cases}$$
(A.2)

Where $c_{m,n}$ is the m-th zero of J_n, $R = 40°$ is the radius of the visual field, and N_{mn} is normalization constant:

$$N_{mn} = \sqrt{\frac{2}{1+\delta_{n,0}}}\,\frac{1}{J_{n+1}(c_{m,n})}$$
(A.3)

A.2. *Symmetric sampling*

In order to calculate the gradient of objective function 18 we use symmetric parameter sampling introduced for parameter exploring policy gradients.[16] We generate J perturbations $\boldsymbol{\epsilon}_j$ from normal distribution $N(\mathbf{0}, \text{diag}(\boldsymbol{\sigma}))$ and create symmetric parameter samples $\boldsymbol{\theta}_j^+ = \boldsymbol{\theta} + \boldsymbol{\epsilon}_j$ and $\boldsymbol{\theta}_j^- = \boldsymbol{\theta} - \boldsymbol{\epsilon}_j$, where $\boldsymbol{\theta}$ is current value of observation parameters. After this we estimate log-probabilities $\log p(\mathcal{X}|\boldsymbol{\theta}_j^+)$ and $\log p(\mathcal{X}|\boldsymbol{\theta}_j^-)$ according to A.3. Using 19 and computing derivatives $\nabla_{\boldsymbol{\rho}} \log p(\boldsymbol{\theta}|\boldsymbol{\rho})$[17] we obtain expression for gradient in the direction of mean parameters μ:

$$\nabla_{\mu_i} \log p(\mathcal{X}|\boldsymbol{\rho}) = \frac{\sum_{j=1}^J \epsilon_{i,j}\big(\log p(\mathcal{X}|\boldsymbol{\theta}_j^+) - \log p(\mathcal{X}|\boldsymbol{\theta}_j^-)\big)}{2J\sigma_i^2}.$$
(A.4)

For estimation of gradient in the direction of standard deviations σ_i, we compare the mean of two log-probabilities with baseline value of log-probability Λ averaged over several computational epochs:

$$\nabla_{\sigma_i} \log p(\mathcal{X}|\boldsymbol{\rho}) =$$
$$\left(\frac{\sum_{j=1}^J \log\big(p(\mathcal{X}|\boldsymbol{\theta}_j^+)\,p(\mathcal{X}|\boldsymbol{\theta}_j^-)\big)}{2J} - \Lambda\right)\left(\frac{\epsilon^2 - \sigma_i^2}{\sigma_i}\right)$$
(A.5)

These expressions for gradient are used in computational scheme for update of parameters $\boldsymbol{\rho}$.

A.3. *Probability estimation*

In this subsection we present mathematical rationale for inference algorithm.

A.3.1. *Probability of history*

In the Infomax model decision where to fixate next is:

$$\pi: \quad \overline{\mathbf{x}}_{k+1} \leftarrow \arg\max_{\overline{\mathbf{x}}} E\big[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})\big].$$
(A.6)

The term $E\big[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})\big]$ is related to the visibility map by[12]

$$E\big[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})\big] = \frac{1}{2}\sum_{\ell=1}^L b_{k,\ell}\, V^2(\overline{\mathbf{x}}, \mathbf{u}_\ell)$$
(A.7)

Due to stochasticity of observation input, the belief state b_k and expected information gain $E\big[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})\big]$

are random variables. Our method of estimation of probability is based on assumption of multivariate Normal distribution of expected information gain, which was verified by implementation of tests for multivariate Normality at section A.3.2. For brevity:

$$\Delta_k = E\left[\Delta I(\mathbf{b}_k, \overline{\mathbf{x}})\right] \qquad (A.8)$$

The probability of the next step in the history $\mathcal{X}$ depends on all previous ones and parameters of visibility map $\boldsymbol{\theta}$:

$$p(\mathcal{X}|\boldsymbol{\rho}) = \prod_k p(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \ldots, \mathbf{x}_0, \boldsymbol{\theta}) \qquad (A.9)$$

The probability of history may be estimated as an integral:

$$p(\mathcal{X}|\boldsymbol{\theta}) = \prod_k \int_{\Psi} p(\mathbf{x}_{k+1}|\Delta_k)\, p(\Delta_k, \mathbf{x}_k, \ldots, \mathbf{x}_0, \boldsymbol{\theta})\, \mathrm{d}\Delta_k$$
$$(A.10)$$

The term $p(\mathbf{x}_{k+1}|\Delta_k)$ is related to the saccade length, according to

$$p(\mathbf{x}_{k+1}|\Delta_k) \propto \exp\left(-\frac{|\mathbf{x}_{k+1} - \overline{\mathbf{x}}_k|^2}{2\zeta_0^2}\right)$$

Where $\overline{\mathbf{x}}_k = \arg\max_{\overline{\mathbf{x}}} \Delta_k$ and the standard deviation ζ_0 from 8. In order to avoid singularity of $\log(p(\mathcal{X}|\boldsymbol{\theta}))$ in the case if probability tends to zero, we use the following regularization:

$$\log\big(p(\mathcal{X}|\boldsymbol{\theta})\big) = \sum_k \log\left(\int_{\Psi} p(\mathbf{x}_{k+1}|\Delta_k)\, p(\Delta_k)\, \mathrm{d}\Delta_k + \omega\right)$$
$$(A.11)$$

Where $\omega = 10^{-7}$ is a regularization constant. Using assumptions of multivariate Normality of expected information gain Δ_k, we compute the mean vector and covariance matrix of Δ_k according to equations A.13 and A.14. On the next step, Monte Carlo integration is used to estimate A.11.

A.3.2. *Distribution of expected information gain*

The belief state is evaluated by Bayesian inference,[12] using the expression:

$$b_{k,\ell} = \frac{b_{0,\ell} \exp\left(\sum_k s_{k,\ell} V^2(\mathbf{u}_\ell - \mathbf{x}_k)\right)}{\sum_\ell b_{0,\ell} \exp\left(\sum_k s_{k,\ell} V^2(\mathbf{u}_\ell - \mathbf{x}_k)\right)} \qquad (A.12)$$

We estimate the mean and the variance of the belief state variable $b_{k,\ell}$, using Monte Carlo sampling of the observation variables $s_{k,\ell}$. Our simulations confirmed that the correlation between the components of $b_{k,\ell}$

at different locations is negligible, and asymptotically tends to zero as $1/D$, where D is the total number of components, or the dimensionality of the computational grid.

The summation operation A.8 is implemented on an array of independent and non-identically distributed variables $b_{k,\ell}$, which are the components of $\mathbf{b}_k$ at different locations. According to the Lyapunov Central Limit Theorem, their sum will converge to a Normal distribution at each location. In order to check the joint normality of the Δ_k components, we used several tests for multivariate normality: Small's test on univariate skewness and kurtosis,[36] Mardia's test on multivariate skewness and kurtosis,[37] Szekely-Rizzo[38] and Wang-Hwang.[39] The results are listed in table A.1, below.

Test name	Test statistic	p-value
Small Q1	15.283	0.761
Small Q2	17.211	0.639
Small Q3	28.66	0.909
Mardia Skewness	18.638	0.401
Mardia Kurtosis	437.451	0.336
Szekely-Rizzo	2.35	0.2
Wang-Hwang	0.99	0.2

Table A.1   Results of multivariate normality tests on the Δ_k information gain distributions.

The information gain was sampled using different VMs and eye-movement trajectories. We used 50 configurations of VM, and applied multivariate normality tests for the samples of information gain variable, with size of 1000. The typical values of test statistics and p-values are presented in A.1. The multivariate normality tests did not reject the hypothesis of joint normality of information gain variables, for any VM configuration. The mean vector and covariance matrix of information gain are:

$$E\big[\Delta_k\big] = \frac{1}{2} \sum_{\ell=1}^L E[b_{k,\ell}]\, V^2(\mathbf{u}_\ell - \overline{\mathbf{x}}) \qquad (A.13)$$

$$\Sigma_{i,j}(\Delta_k) = \frac{1}{4} \sum_{\ell=1}^L V^2(\mathbf{u}_\ell - \overline{\mathbf{x}}_i)\, \Sigma_{k,\ell}(\mathbf{b})\, V^2(\mathbf{u}_\ell - \overline{\mathbf{x}}_j)$$
$$(A.14)$$

where $\Sigma_{k,\ell}(\mathbf{b})$ is the covariance of of the belief vector $\mathbf{b}$. These expressions use the classical definitions of error propagation, through linear transformations.

The amplitude and noise in information gain variable define the decision-making behavior of the ob-

server. Because information gain variable is random, the location of its maxima will vary with roll-out. However, locations with the highest amplitude and noise will be most likely to be selected. The locations with the highest amplitude of information gain correspond to unexplored ones. Because, the AMD patient is effectively blind in scotoma locations, it's a reasonable strategy to make the next fixation to the one of scotoma locations. However, the choice of the next location of fixation is also influenced by the noise in information gain, which will correspond to false positive locations of targets on periphery. For this reason, it's not always the case when scotoma directions are the preferable ones, as it was established in section 2.3. We hypothesize that the reason of the difference in behavior between the low and the high visibility cases is corresponding domination of amplitude-based and noise-based selections in these cases.

A.4.   *Influence of length of training set*

It is important that the eye-movement trajectory contains enough information for inference of the observation model. In this section we investigate the influence of length of training set on performance of the inference algorithm.

A set of 10 ground truth visibility maps were generated by random selection of each observation parameter $\boldsymbol{\theta}$, as defined in A.1, over the interval $[-1, 1]$. For each visibility map function, we generate eye-movement trajectories with lengths $[12, 24, 48, 96, 192]$, and run episodes of Po-Mdp on a 64×64 computational grid. These generated trajectories served as an input for the inference algorithm.

Figure A.1 shows the Euclidean distance between ground truth and inferred VMs, on the final epoch, according to the length of the training set. As expected, both the variance and mean of the distance fall, as the number of episodes in the training set increases. The performance of the algorithm stabilizes after a certain number of episodes, because estimation of the exact VM is impossible, due to uncertainty of saccade execution and noisy estimation of the probability integral. In summary, we recommend using data sets that contain at least as many fixations as there are cells in the computational grid.
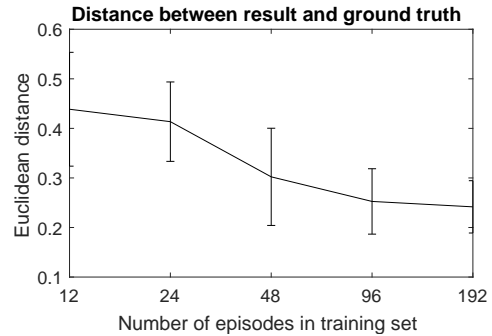

Fig. A.1.   Euclidean distance between ground truth and inferred visibility map on final epoch depending on the number of episodes.

References

1. E. Wiecek, M. L. Jackson, S. C. Dakin, and P. Bex, "Visual search with image modification in age-related macular degenerationvisual search in amd," *Investigative ophthalmology & visual science*, vol. 53, no. 10, pp. 6600–6609, 2012.
2. S. Van der Stigchel, R. A. Bethlehem, B. P. Klein, T. T. Berendschot, T. Nijboer, and S. O. Dumoulin, "Macular degeneration affects eye movement behavior during visual search," *Frontiers in psychology*, vol. 4, p. 579, 2013.
3. C. P. Janssen and P. Verghese, "Training eye movements for visual search in individuals with macular degeneration," *Journal of vision*, vol. 16, no. 15, pp. 29–29, 2016.
4. M. Thibaut, C. Delerue, M. Boucart, and T. H. C. Tran, "Visual exploration of objects and scenes in patients with age-related macular degeneration," *Journal francais d'ophtalmologie*, vol. 39, no. 1, pp. 82–89, 2016.
5. F. Geringswald, A. Herbik, M. B. Hoffmann, and S. Pollmann, "Contextual cueing impairment in patients with age-related macular degeneration," *Journal of Vision*, vol. 13, no. 3, pp. 28–28, 2013.
6. 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, 2013.
7. M. P. Eckstein, "Visual search: A retrospective," *Journal of vision*, vol. 11, no. 5, pp. 14–14, 2011.
8. N. J. Butko and J. R. Movellan, "Infomax control of eye movements," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 2, pp. 91–107, 2010.
9. J. Najemnik and W. S. Geisler, "Optimal eye movement strategies in visual search," *Nature*, vol. 434, no. 7031, pp. 387–391, 2005.
10. M. P. Eckstein, W. Schoonveld, S. Zhang, S. C. Mack, and E. Akbas, "Optimal and human eye movements to clustered low value cues to increase decision rewards during search," *Vision research*, vol. 113, pp. 137–154, 2015.
11. J. Najemnik and W. S. Geisler, "Simple summation

rule for optimal fixation selection in visual search," *Vision Research*, vol. 49, no. 10, pp. 1286–1294, 2009.

12. J. Najemnik and W. S. Geisler, "Eye movement statistics in humans are consistent with an optimal search strategy," *Journal of Vision*, vol. 8, no. 3, pp. 4–4, 2008.

13. A. Y. Vasilyev, "Optimal control of eye-movements during visual search," *IEEE Transactions on Cognitive and Developmental Systems*, 2018.

14. Y. Tsank and M. P. Eckstein, "Domain specificity of oculomotor learning after changes in sensory processing," *Journal of Neuroscience*, vol. 37, no. 47, pp. 11469–11484, 2017.

15. D. V. Walsh and L. Liu, "Adaptation to a simulated central scotoma during visual search training," *Vision Research*, vol. 96, pp. 75–86, 2014.

16. F. Sehnke, C. Osendorfer, T. Rückstieß, A. Graves, J. Peters, and J. Schmidhuber, "Parameter-exploring policy gradients," *Neural Networks*, vol. 23, no. 4, pp. 551–559, 2010.

17. F. Sehnke, A. Graves, C. Osendorfer, and J. Schmidhuber, "Multimodal parameter-exploring policy gradients," in *Machine Learning and Applications (ICMLA), 2010 Ninth International Conference on*, pp. 113–118, IEEE, 2010.

18. T. Rückstieß, M. Felder, and J. Schmidhuber, "State-dependent exploration for policy gradient methods," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 234–249, Springer, 2008.

19. T. Rückstiess, F. Sehnke, T. Schaul, D. Wierstra, Y. Sun, and J. Schmidhuber, "Exploring parameter space in reinforcement learning," *Paladyn*, vol. 1, no. 1, pp. 14–24, 2010.

20. R. Jaulmes, J. Pineau, and D. Precup, "Active learning in partially observable markov decision processes," in *ECML*, pp. 601–608, Springer, 2005.

21. R. Jaulmes, J. Pineau, and D. Precup, "Learning in non-stationary partially observable markov decision processes," in *ECML Workshop on Reinforcement Learning in non-stationary environments*, vol. 25, pp. 26–32, 2005.

22. A. Atrash and J. Pineau, "A bayesian method for learning pomdp observation parameters for robot interaction management systems," in *ICAPS POMDP Practitioners Workshop*, 2010.

23. N. J. Butko and J. R. Movellan, "I-pomdp: An infomax model of eye movement," in *2008 7th IEEE International Conference on Development and Learning*, pp. 139–144, IEEE, 2008.

24. A. Y. Vasilyev, "Vm.inference." https://github.com/alexander-vasilyev/VM.Inference, 2018.

25. A. M. McKendrick, "Recent developments in perimetry: test stimuli and procedures," *Clinical and Experimental Optometry*, vol. 88, no. 2, pp. 73–80, 2005.

26. D. J. Taylor, N. D. Smith, and D. P. Crabb, "Searching for objects in everyday scenes: Measuring performance in people with dry age-related macular degeneration," *Investigative ophthalmology & visual science*, vol. 58,

no. 3, pp. 1887–1892, 2017.

27. L. McIlreavy, J. Fiser, and P. J. Bex, "Impact of simulated central scotomas on visual search in natural scenes," *Optometry and vision science: official publication of the American Academy of Optometry*, vol. 89, no. 9, p. 1385, 2012.

28. J. A. Jacko, A. B. Barreto, J. Y. Chu, I. U. Scott, R. H. Rosa Jr, and C. C. Pappas Jr, "Macular degeneration and visual search: What we can learn from eye movement analysis," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 44, pp. 116–119, SAGE Publications Sage CA: Los Angeles, CA, 2000.

29. J. A. Jacko, I. U. Scott, A. B. Barreto, H. S. Bautsch, J. Chu, and W. B. Fain, "Iconic visual search strategies: A comparison of computer users with amd versus computer users with normal vision," in *Proceedings of the Ninth International Conference on Human–Computer Interaction*, vol. 2, p. 423, 2001.

30. W. Seiple, J. P. Szlyk, T. McMahon, J. Pulido, and G. A. Fishman, "Eye-movement training for reading in patients with age-related macular degeneration," *Investigative Ophthalmology & Visual Science*, vol. 46, no. 8, pp. 2886–2896, 2005.

31. R. Engbert, A. Nuthmann, E. M. Richter, and R. Kliegl, "Swift: a dynamical model of saccade generation during reading.," *Psychological review*, vol. 112, no. 4, p. 777, 2005.

32. S. K. Mannan, A. L. Pambakian, and C. Kennard, "Compensatory strategies following visual search training in patients with homonymous hemianopia: an eye movement study," *Journal of neurology*, vol. 257, no. 11, pp. 1812–1821, 2010.

33. B. Fuglede and F. Topsoe, "Jensen-shannon divergence and hilbert space embedding," in *Information Theory, 2004. ISIT 2004. Proceedings. International Symposium on*, p. 31, IEEE, 2004.

34. B. Efron and C. Stein, "The jackknife estimate of variance," *The Annals of Statistics*, vol. 9, no. 3, pp. 586–596, 1981.

35. V. Lakshminarayanan and A. Fleck, "Zernike polynomials: a guide," *Journal of Modern Optics*, vol. 58, no. 7, pp. 545–561, 2011.

36. N. Small, "Marginal skewness and kurtosis in testing multivariate normality," *Applied Statistics*, pp. 85–87, 1980.

37. K. V. Mardia, "Measures of multivariate skewness and kurtosis with applications," *Biometrika*, vol. 57, no. 3, pp. 519–530, 1970.

38. N. K. Bakirov, M. L. Rizzo, and G. J. Székely, "A multivariate nonparametric test of independence," *Journal of Multivariate Analysis*, vol. 97, no. 8, pp. 1742–1756, 2006.

39. C.-C. Wang and Y.-T. Hwang, "A new functional statistic for multivariate normality," *Statistics and Computing*, vol. 21, no. 4, pp. 501–509, 2011.

Appendix B (supplementary material)
Direct optimization of the objective function

In this section we discuss alternative approaches for solving optimization problem 16, for comparison with the gradient-based method developed in section 3 of the main text. The principal challenges in this problem are:

- Noisy estimation of log-probability due to Monte Carlo integration.
- Presence of many local maxima in the multidimensional parameter space

We selected three standard algorithms for comparison with the gradient method: Nelder-Mead simplex direct search, derivative-free Pattern Search, and Simulated Annealing, which are implemented in Matlab Global Optimization Toolbox.

The set of 10 ground truth visibility maps were generated by random selection of each observation parameter $\boldsymbol{\theta}$ A.1 in the interval $[-1, 1]$. For each VM function, we generate eye-movement trajectories with lengths 192, and run episodes of Po-Mdp on a 64×64 computational grid. The value of the objective function was evaluated according to equation A.11. The results of each algorithm should be compared with those for the gradient algorithm (Fig. 7). Termination criteria for all three algorithms was a lower bound on step size $|\Delta\theta| < 0.02$. The initial mesh size for the Pattern Search algorithm was set to 0.2 for these experiments. The initial temperature parameter for Simulated Annealing was set to $\lambda = 100$. The average distance to the target VM, for all three algorithms, was close to that of the gradient method, $E[R] = 0.2$, however, the standard deviations of the distances exceeded those of the gradient, $\sigma[R] = 0.05$. Overall, the Pattern Search algorithm performed better than the alternatives. We preferred using gradient method due to lower mean and standard deviation of distances comparing to alterna-tives, which means higher precision and reliability in solving the optimization problem. Figures B.1, B.2, and B.3 show the performance of the three different algorithms.
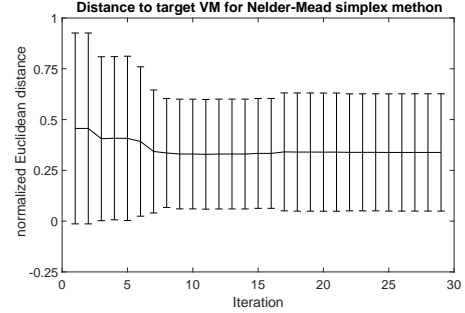

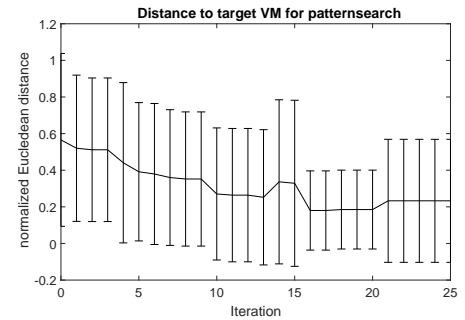Fig. B.1.   Performance of Nelder-Mead simplex algorithm.
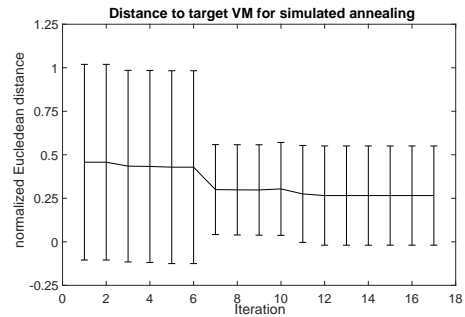

Fig. B.2.   Performance of the Pattern Search algorithm.

Fig. B.3.   Performance of the Simulated Annealing algorithm.