

Parameter Recovery of a Visual-Vestibular Interactions

Model using Adaptive Stimulus Selection

Bachelor's Thesis in Artificial Intelligence

Supervisor: L. P. J. Selen; Artificial Intelligence, Donders Institute
Radboud University

Niels van Nistelrooij (s4713648)

June 28, 2019

Abstract

The signal quality of the vestibular organ diminishes with age. This can have negative effects on balance, coordination and vision. The aim of the project was to develop techniques that radically decrease the number of trials needed in the Rod-and-Frame task for testing vestibular functionality based on a Bayesian model. The model is a reduced version of the model by Alberts et al. (2016) where the head orientation is fixed at upright. The mean inferred parameter values in that paper were used to make a generative model that generates responses given stimuli. Adaptive stimulus selection was used to select the next stimulus, such that the corresponding response was maximally informative for the model parameter recovery. This was done in silico, by adaptively manipulating the rod and frame orientations. The proposed model's parameters could not be successfully recovered. However, the advantage of adaptive stimulus selection over random stimulus selection was effectively shown when recovering single parameter and parameter pair values. This advantage was not observed for the model with all free parameters, suggesting that the proposed model's parameters are fundamentally unrecoverable. The discussion provides a number of suggestions on how this unrecoverability might be overcome.

1 Introduction

The vestibular organ in the inner ear facilitates inertial perception and spatial orientation [6]. To do this, it uses its semicircular canals to provide information about the rotational velocity of the head [3]. The otolith organs, on the other hand, give information about the gravito-inertial force, which is the sum of the gravitational force and the inertial force from linear acceleration [7]. So the otolith organs give information about the orientation of the head relative to gravity and linear acceleration. These three functions combined, make it, so that the vestibular organ is involved in a large number of neural operations. One such operation is determining the orientation of the head in space relative to gravity. Besides the otolith organs, visual cues are used and bias the orientation perception toward the cardinal directions of the visual object [10].

To quantify the bias, the Rod-and-Frame task is used [4]. The subjects look at a screen that presents a briefly flashed rod inside a square frame. The rod and frame have different orientations relative to gravity. The subject then judges whether the rod is rotated clockwise or counterclockwise relative to gravity and presses a corresponding button. The data of this experiment reliably finds that subjects perceive the rod's orientation with a bias away from the orientation of the frame. If the frame is, for example, at 10° clockwise and the rod is at 0° , the subject's rate for pressing the clockwise button is significantly lower than 50%. Put differently, the subject's head orientation perception is biased to the clockwise direction following the orientation of the frame and the upright rod then seems counterclockwise in comparison. Furthermore, if the subject's head is oriented 30° right ear down, they receive more noisy information from their otolith organs about the head orientation. This in turn causes the visual bias to become greater, because the subject is relying more on the visual information than the information from the otolith organs [1].

Another cause for a greater bias is a progressive autoimmune disease, DFNA9 [2]. Patients experience a much faster decline of their hearing and vestibular functioning compared to healthy subjects and completely lose their hearing and vestibular organs in later stages of life. Although their hearing can be partially brought back with the use of cochlear implants, too little is known about the vestibular organs to apply a similar solution to bring back vestibular functioning. To cope with their loss of vestibular information, patients more strongly use visual cues in their environment to keep balance and orientation perception. This means that, when these patients perform the Rod-and-Frame task, the data shows a bigger bias away from the orientation of the frame [9]. Patients can use non-vestibular information, such as neck proprioceptive information, to get indirect information about the orientation of the head relative to gravity. Both patients and controls weigh their vestibular/non-vestibular information similarly, so the difference in signal precision between patients and controls only becomes apparent in the visual integration. The patients rely more on the visual information and will weigh that information higher than the controls. So the resulting increase in visual bias for the patients is due to an increase in visual information weight compared to controls and not a decrease of the

vestibular/non-vestibular information weight.

To explain these data, a Bayesian model has been devised that captures vestibular, visual and prior information and tries to explain the visual bias [1]. In the present thesis, a reduced version of this model will be used. The number of parameters will be decreased, making the model simpler and inference possible on consumer hardware. Inferring the model parameter values in the original experiment took 1620 trials per subject per head orientation. The rods and frames that were presented to the subjects were chosen pseudo-randomly, such that each rod-frame combination would occur exactly 10 times for each head orientation. In the present paper, only 500 iterations will be used to recover the model parameters using adaptive stimulus selection. Instead of pseudo-randomly selecting the next stimulus, the next stimulus is computed to be the most informative for determining the parameters of the model [11]. This is done, in order to reduce the time to complete an empirical experiment, while still having comparable model parameter recovery performance. The inferred model parameter values can then be used to track the decline in vestibular functioning in elderly people or DFNA9 patients. Reducing the time for such an experiment to a clinically applicable time, means that the progression of the decline of vestibular functioning could be tracked in practice.

The main focus of the present thesis is to successfully recover the parameter values of a modified version of the model in [1]. Parameter recovery is done using random and adaptive stimulus selection, thereby comparing the two methods of selecting stimuli. First, the modified visual-vestibular interactions model will be introduced. Then generative models, which make use of the mean inferred parameter values in [1] and [9], are presented, which give responses for given stimuli. Model parameter recovery and adaptive stimulus selection are explained to get a better understanding of the idea and maths behind it. Then the results are presented, where, first, single parameter and parameter pairs are tried to be recovered, before the model with all free parameters is tried to be recovered. The thesis is concluded with a discussion and conclusion.

2 Methods

2.1 Visual-Vestibular Interactions Model

For the present project, the model from Alberts et al. (2016) [1] served as the starting point. The Bayesian model consists of a head-in-space prior, a vestibular likelihood and a visual contextual likelihood that are combined to form the posterior head-in-space orientation estimate. The prior and vestibular likelihood distributions are Gaussians with a particular σ_{HP} , respectively σ_{HS} , where the prior is centered at 0° and the vestibular likelihood is centered at the orientation of the subject's head relative to gravity. The visual contextual likelihood consists of a sum of four von Mises distributions with peaks at the cardinal directions of the stimulus frame. The posterior is then retrieved by multiplying the three distributions and normalizing the result. Rod, frame and

head orientations were used as a stimulus each trial. The head orientation determines the vestibular likelihood, the frame orientation determines the visual contextual likelihood and the rod orientation determines the rod-in-space orientation estimate given the posterior head-in-space orientation estimate. This rod-in-space orientation estimate is in turn used to make the decision between responding clockwise or counterclockwise.

For the current paper, the head orientation is excluded from the stimulus. Instead, the orientation of the subject's head will be fixed at upright relative to gravity. Figure 1 shows the proposed vestibular model and shows examples of the likelihoods and the posterior distribution.

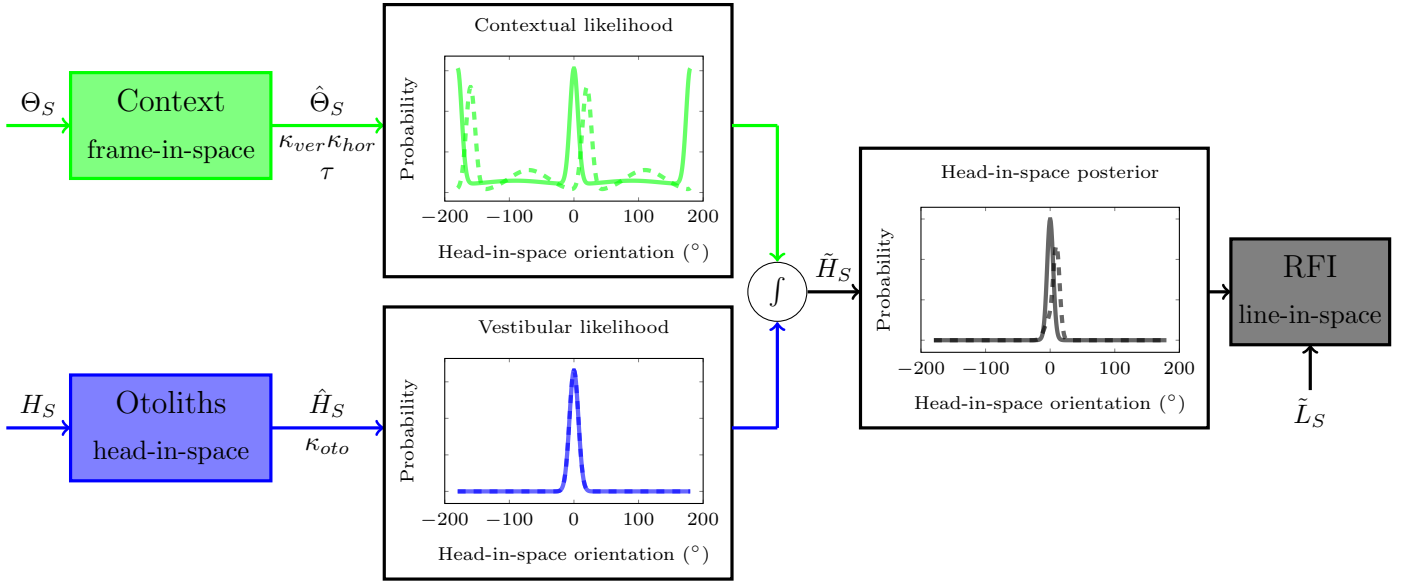


Figure 1: Graphic representation of the reduced Bayesian visual-vestibular interactions model. Physical signals about the frame-in-space orientation (Θ_S) and head-in-space orientation (H_S) are transformed into sensory signals, denoted by the hat symbol ($\hat{\cdot}$). Sensory signals are assumed to be accurate up to a certain precision ($\kappa_{ver}, \kappa_{hor}$ and κ_{oto} , respectively). For an optimal estimate of head-in-space orientation, denoted by a tilde ($\tilde{\cdot}$), the model combines the visual contextual likelihood $P(\hat{\Theta}_S|H_S)$ together with the vestibular likelihood $P(\hat{H}_S|H_S)$. To achieve this, the individual probability distributions are multiplied: $P(\tilde{H}_S|\hat{H}_S, \hat{\Theta}_S) = P(\hat{H}_S|H_S) \cdot P(\hat{\Theta}_S|H_S)$. The maximum of the resulting posterior distribution (MAP) is selected as the perceived head-in-space orientation (\tilde{H}_S), whereas the width of the posterior is a measure of the response variability. The perceived orientation of the line in space is then obtained by a coordinate transformation using the rod-in-space orientation (\tilde{L}_S). The probability distributions in the figure represent the case in which the subject is presented with a frame displayed upright (solid, $\Theta_S = 0^\circ$) and a frame oriented right relative to gravity (dashed, $\Theta_S = 20^\circ$). Figure adapted from Figure 2 in [1].

2.1.1 Visual Contextual Likelihood

The visual contextual likelihood for a given frame-in-space orientation (Θ_S) is given by:

$$P(\hat{\Theta}_S | H_S, \Theta_S, \kappa_{ver}, \kappa_{hor}, \tau) = \sum_{i=1}^4 \frac{\exp\{\kappa(i) \cdot \cos[\varphi(i) - \Theta_S - H_S]\}}{2\pi \cdot I_0[\kappa(i)]}$$

where $\kappa(i) = [\kappa_1, \kappa_2, \kappa_1, \kappa_2]$,
 $\kappa_1 = \kappa_{ver} - [1 - \cos(|2 \cdot \Theta_S|)] \cdot \tau \cdot (\kappa_{ver} - \kappa_{hor})$,
 $\kappa_2 = \kappa_{hor} + [1 - \cos(|2 \cdot \Theta_S|)] \cdot (1 - \tau) \cdot (\kappa_{ver} - \kappa_{hor})$,
 $\varphi(i) = [0^\circ, 90^\circ, 180^\circ, 270^\circ]$
and I_0 is the modified Bessel function of the first kind.

An example of such a distribution can be seen in Figure 1. It consists of four von Mises distributions with their peaks at 0° , 90° , 180° and 270° plus the frame's orientation. The difference between the vertical and horizontal peaks is dependent on κ_{ver} and κ_{hor} and how the peaks change as the frame rotates in space, is determined by τ . There is a clear distinction between the vertical and horizontal cardinal directions' effect on the perceived head-in-space orientation, where the vertical cardinal direction has a stronger effect on the perceived head-in-space orientation than the horizontal direction. If the frame is rotated 45° in space, $\kappa_1 = \kappa_2$, so the vertical and horizontal peaks of the distribution are equal. The visual contextual likelihood still reduces uncertainty, but does not introduce bias at this frame orientation. The κ parameter can be interpreted as the precision of the distribution. The relation between the variance and precision is $\kappa \approx \frac{1}{\sigma^2}$, so the κ parameter can be seen as the inverse of the variance.

2.1.2 Vestibular Likelihood

The vestibular likelihood is also a von Mises distribution given by:

$$P(\hat{H}_S | H_S, \kappa_{oto}) = \frac{\exp\{\kappa_{oto} \cdot \cos[-H_S]\}}{2\pi \cdot I_0[\kappa_{oto}]}$$

An example can be seen in Figure 1. The vestibular likelihood is centered around 0° with usually a higher precision, i.e. higher κ parameter, than the visual contextual likelihood. In the model by Alberts et al. (2016) [1], excluding the head orientation from the stimulus causes the vestibular likelihood and head-in-space prior to become both centered at 0° . The vestibular likelihood and head-in-space prior thus become non-dissociable, so these two Gaussian distributions are replaced by a single von Mises distribution centered at 0° . This choice was made, because the von Mises distribution is the circular distribution most similar to a Gaussian distribution. The vestibular likelihood should be a circular distribution, because its input, the head-in-space orientation, is circular. The two variability parameters of the prior and vestibular likelihood (σ_{HP} , respectively σ_{HS}) are then also replaced by a single parameter, κ_{oto} .

2.1.3 Head-in-Space Posterior

The head-in-space posterior probability for a given frame-in-space orientation (Θ_S) is then given by:

$$P(\tilde{H}_S|\hat{H}_S, \hat{\Theta}_S, \Theta_S, \kappa_{ver}, \kappa_{hor}, \tau, \kappa_{oto}) \propto P(\hat{\Theta}_S|H_S, \Theta_S, \kappa_{ver}, \kappa_{hor}, \tau) \cdot P(\hat{H}_S|H_S, \kappa_{oto})$$

See Figure 1 for an example of the posterior head-in-space distribution. After multiplying the likelihoods, the result is normalized to make a probability distribution. The maximum a posteriori (MAP) of the head-in-space posterior is taken as the perceived head-in-space orientation. To then compute the perceived rod-in-space orientation, one subtracts the perceived head-in-space orientation from the line-in-space orientation. Then based on how big the perceived rod-in-space orientation is, the subject’s responses are generated. The width of the head-in-space posterior gives a measure of the response variability.

All in all, the new reduced visual-vestibular interactions model has 4 parameters: κ_{ver} , κ_{hor} and τ for the visual contextual information and κ_{oto} for the vestibular information. κ_{ver} codes the precision of the vertical visual context, whereas κ_{hor} codes the precision of the horizontal visual context, both when the frame is upright. τ modifies how the vertical and horizontal visual contextual precisions change as the frame rotates. Finally, κ_{oto} codes the precision of the vestibular likelihood.

2.2 Experimental Stimuli

The experimental stimuli that will be presented are 9 rods -7° , -4° , -2° , -1° , 0° , 1° , 2° , 4° and 7° relative to gravity and 18 frames from -45° to 40° with steps of 5° . The stimuli are a good compromise between too few, where the choices of the adaptive stimulus selection procedure would be crude, and too many, where the computational resources that the algorithm requires, exceed clinical applicability. Rods with a greater orientation away from the gravity line than $\pm 7^\circ$ would almost always be correctly identified, so going further than $\pm 7^\circ$ does not help parameter recovery.

2.3 Generative Models

Three generative models will be presented. These models will be used to generate clockwise or counterclockwise responses given rod and frame stimuli. The mean inferred parameters in [1] will be used for the young generative model and the mean inferred parameters of the controls and patients in [9] will be used for the old, respectively patient generative models. The population sample in [1] was significantly younger than the control and patient population samples in [9]. The parameter values can be seen in Table 1.




Figure 2: The expected response variability, weights and bias for the generative models as a function of frame orientation. See Table 1 for the used parameter values.

The vestibular likelihood standard deviation is denoted by β_{HS} in [1] and [9]. This parameter is translated to a κ precision and used as κ_{oto} for the present experiment. This way, the vestibular likelihood standard deviation is equal to the mean vestibular likelihood standard deviation in [1] and [9] when the head is not rotated. So the influence of the head-in-space prior is assumed to be negligible when the head is not rotated and therefore κ_{oto} is only determined by the inferred vestibular likelihood standard deviation.

Model	κ_{ver}	κ_{hor}	τ	κ_{oto}	λ
Young	86.24	1.451	0.80	145.3	0.02
Old	45.37	2.552	0.97	71.76	0.02
Patient	84.09	0.8721	0.87	69.29	0.05

Table 1: Parameter values for the generative models.

The response variability, weights and bias of the generative models can be seen in Figure 2. The bias is equal to the point of subjective equivalence; the rod orientation where the probability of responding clockwise is 50%. The response variabilities are taken from the vestibular likelihood and the vertical peaks of the visual contextual likelihood. The weights are computed as $\frac{1/\sigma_i^2}{\sum_j 1/\sigma_j^2}$ [8]. The noise from the vestibular/non-vestibular information is comparable between the old and patient generative models, but the noise from the visual context is much smaller for the patient compared to the old generative model. This will cause a greater bias for a frame orientation around $\pm 20^\circ$, which can indeed be seen in the bottom bias plots.

In the old generative model, the response variability and weight of the otolith organs never intersects with the response variability and weight of the visual context. This indicates that the subject will always weigh the information from the otolith organs higher than the visual information. The bias is still measurably big with a maximum bias of $\pm 5.4^\circ$.

The patient generative model shows a different pattern. The response variability and weight of the vestibular/non-vestibular information do intersect with the visual context. So, for certain frame orientations, the subject will weigh the information from the visual context higher than the vestibular/non-vestibular information. This can also be seen by the bias which is larger than the bias of the old generative model with a maximum of $\pm 8.4^\circ$.

The frame orientations where the maximum of the bias is experienced is $\pm 20.2^\circ$ and $\pm 22.1^\circ$ for the old, respectively patient generative models.

Comparing the young generative model on the one hand and the old and patient generative models on the other hand, it becomes apparent that the vestibular likelihood standard deviation for the young generative model is smaller compared to the old and patient models. This can be explained by the age difference in the population sample of [1] and [9]. The young generative model was made from subjects which were 27 ± 5 years old, whereas the old and patient generative models were made from subjects which were 58.4 ± 10.5 , respectively 65.1 ± 8.0 years old. So the data of these generative models follows the proposition that vestibular functionality diminishes with age, as also highlighted in [8].

2.4 Model Parameter Recovery

The next aim is to see whether the parameters of the generative models can be recovered.

2.4.1 Likelihood

First, a likelihood look-up table is computed using the formulas of the internals of the vestibular model discussed in Section 2.1.

The probability of responding clockwise for a given line-in-space (L_S) and frame-in-space (Θ_S) orientation is given by:

$$P(CW|x, \theta) = \lambda + (1 - 2\lambda) \cdot \int_{-\pi}^{L_S} P(\tilde{H}_S | \hat{H}_S, \hat{\Theta}_S, \Theta_S, \kappa_{ver}, \kappa_{hor}, \tau, \kappa_{oto}) d\tilde{H}_S$$

where $x = [L_S, \Theta_S]$
and $\theta = [\kappa_{ver}, \kappa_{hor}, \tau, \kappa_{oto}, \lambda]$.

With probability $P(CCW|x, \theta) = 1 - P(CW|x, \theta)$, the subject responses counterclockwise.

The lapse rate, λ , is only added in this step, because it is not part of the vestibular model. If the subject lapses, the vestibular model is not used to generate a response. Instead, a random response is generated. The lapse rate must still be mentioned, because it is part of the Rod-and-Frame task responses model and that model accounts for subject lapses. The number of parameters that then need to be recovered is equal to 5 ($\kappa_{ver}, \kappa_{hor}, \tau, \kappa_{oto}, \lambda$).

$P(CW|x, \theta)$ is computed for each combination of parameters and stimuli and saved before the experiments in a big look-up table. Computing the likelihood look-up table takes, by quite a big margin, the longest time. So being able to save the look-up table before the experiments greatly decreases the time it takes to run the experiments.

The difference between the inference program's likelihood look-up table and the generative agent's, is that the generative agent only has one combination of parameters, which reduces the likelihood look-up table for the generative agent to $P(CW|x)$.

Instead of using the set of rods denoted in Section 2.2, the cumulative posterior is computed for 10,000 rods from $-\pi$ to π . Then using spline interpolation, the cumulative posterior is reduced to the 9 rods denoted in Section 2.2. Finally, the λ parameter is added to the cumulative posterior to compute the likelihood distribution.

The cumulative posterior needs to be made with the full input, i.e. L_S from $-\pi$ to π , which could not be done with the rods in Section 2.2. 10,000 was chosen to ensure that the error when reducing is minimal. The final likelihood distribution only uses the stimuli rods, so the experiments will be a lot faster. Indeed, computing the cumulative posterior for 10,000 rods is the reason that computing the likelihood look-up table takes such a large amount of time.

Previously, a nearest-neighbour technique was used to reduce the cumulative posterior; instead of using spline interpolation, the rods that were closest to the rods in Section 2.2 were used for the likelihood distribution. This introduced small errors, as the chosen rods were not exactly the stimuli rods and the model parameter recovery became worse at later trials. With spline interpolation, the cumulative posterior is fitted to a number of continuous cubic functions.

The cumulative posterior of the stimuli rods can then be exactly gathered from these continuous cubic functions. The small error was reduced by using spline interpolation and the performance of model parameter recovery improved in later trials.

2.4.2 Prior Probability over Parameter Combinations

First, the prior probability is chosen for each parameter on its own: $P(\kappa_{ver})$, $P(\kappa_{hor})$, $P(\tau)$, $P(\kappa_{oto})$, $P(\lambda)$. τ and λ use beta distributions and the other three parameters use uniform distributions. Both the τ and λ parameters can have values from 0 to 1 and looking at [1] and [9], it seems likely that these parameters follow a beta distribution in the population samples. The prior parameter distributions can be seen in Figure 3. To still be able to converge to a τ of 1.0 or a λ of 0.0, the prior probability density for extreme parameter values is increased to a threshold. Otherwise, the prior probability for $\tau = 1.0$ and $\lambda = 0.0$ would be equal to 0.0, meaning that extreme values such as these would be impossible to recover.

The threshold was chosen, such that the ratio between the minimum and maximum prior probability densities was equal to 10. See Table 2 for the parameter prior shape parameters and parameter thresholds for τ and λ . The final parameter distribution precisions for τ and λ increased when the priors were introduced. Because of the priors, the distributions of τ and λ had a higher precision at the start of each experiment compared to the κ parameters. This caused the precision at the end of the experiment to also increase.

The other parameters have uniform prior distributions, because no clear distribution could be ascertained from the parameter values in [1] and [9].




Figure 3: Probability density functions of priors of τ and λ . See Table 2 for the beta distribution shape parameters that were used. The density thresholds of τ and λ are equal to 0.385, respectively 0.632.

Parameter	τ	λ
α	10	2.0
β	1.6	35
Threshold	0.647	0.132

Table 2: Beta distribution shape parameters of priors of τ and λ and the corresponding parameter value thresholds.

Then, the prior probability for each parameter combination is computed by first computing the Cartesian product of all the parameter prior probabilities and then taking the product of each element:

$$P(\theta) := \left\{ \prod p \mid p \in P(\kappa_{ver}) \times P(\kappa_{hor}) \times P(\tau) \times P(\kappa_{oto}) \times P(\lambda) \right\}$$

So $P(\theta)$ will be a long vector where each element gives a prior probability for a parameter combination. Putting the prior probabilities in this configuration makes inference computations faster.

2.4.3 Inference

The inferred posterior parameter distribution is given by [11]:

$$P(\theta|D_t) \propto P(\theta) \cdot \prod_{s=1}^t P(y_s|x_s, \theta)$$

where $D_t := \{x_s, y_s\}_{s=1}^t$,
 $x = [L_S, \Theta_S]$
and $y := \{CW, CCW\}$.

The \propto symbol denotes that the posterior distribution needs to be normalized after each multiplication.

At the start of the experiment, $P(\theta|x, y) \propto P(\theta) \cdot P(y|x, \theta)$ is computed. Adaptive stimulus selection uses this posterior distribution to select the first stimulus, whereas random stimulus selection picks a random first stimulus. The stimulus is given to the generative model, which will generate a clockwise or counterclockwise response. This first stimulus-response pair (x_1, y_1) is used to select $P(\theta|x_1, y_1)$ from the posterior, which is used as prior in the posterior update: $P(\theta|x, y) \propto P(\theta|x_1, y_1) \cdot P(y|x, \theta)$. Then for each iteration of the experiment, the stimulus is selected using the updated posterior or it is selected randomly, the response is generated by the generative model and the posterior is updated with $P(\theta|x, y) \propto P(\theta|x_s, y_s) \cdot P(y|x, \theta)$.

At the end of the experiment, the posterior parameter combination distribution is $P(\theta|D_t) = P(\theta|x_t, y_t)$. This distribution has a probability for each combination of parameters. The highest probability, or the MAP value, of this distribution gives the combination of parameters that are most likely, given the

stimulus-response pairs. The mean value of the posterior distribution can also be computed by doing a matrix multiplication between the Cartesian product of the parameter values and the posterior parameter combination distribution, i.e. $(\kappa_{ver} \times \kappa_{hor} \times \tau \times \kappa_{oto} \times \lambda) \cdot P(\theta|D_t)$.

The posterior is normalized after it has been initialized and after it has been updated each trial. If this was not done, the probabilities of the distribution would quickly decrease to 0.0, because of the product and the finite precision of numbers in computer programs.

Updating the posterior iteratively allows analysis computations each trial using the intermediate posterior. The results of these computations will later in the thesis be used to present figures that show the inference throughout the experiment.

2.5 Stimulus Selection

Two methods of selecting stimuli will be presented; random and adaptive stimulus selection. The first selects randomly and the second selects based on a heuristic.

2.5.1 Random Stimulus Selection

Random stimulus selection will randomly select the stimulus for the next trial. The rod and frame are selected independently of each other. No constraints on duplicates, the order or skew are set; it is assumed that the random number generator handles most problems on its own.

The choice could be made to use pseudo-random stimulus selection, where each stimulus appears a fixed number of times during the experiment, but this choice was not made. If this choice would be made, then the number of iterations of the experiment is limited to multiples of $9 \cdot 18 = 162$. This would decrease the flexibility in the choice for the number of iterations of the experiment.

2.5.2 Adaptive Stimulus Selection

For adaptive stimulus selection, the stimulus with minimum expected entropy in the posterior distribution will be selected. Using the minimum of the expected entropy is a good heuristic for selecting the next stimulus, because it is a measure of expected information. In other words, the stimulus with the minimum expected entropy will be the most informative for learning the model parameters and thus fitting the data. The stimulus for the next trial will be computed as follows [11].

First, the marginalized distribution of each stimulus (x) and response (y) combination is computed as:

$$P(x_i, y_j) = \int P(\theta|x_i, y_j)d\theta$$

$P(x, y)$ is also used in the inference to normalize the posterior. So the marginalized distribution should be computed before normalizing the posterior. Next, the posterior is normalized and then the entropy is computed for each stimulus-response combination as:

$$H(x_i, y_j) = - \int P(\theta|x_i, y_j) \log P(\theta|x_i, y_j) d\theta$$

Then the expected entropy is computed as:

$$\mathbb{E}[H(x)] = \sum_{i=1}^2 P(x, y_i) \cdot H(x, y_i)$$

Lastly, to select the next stimulus, the minimum of the expected entropy is picked:

$$x_{t+1} = \arg \min_x \mathbb{E}[H(x)]$$

2.6 Experiments

Before a set of experiments, the likelihood look-up table, $P(CW|x, \theta)$, is computed and saved. The generative agent’s look-up table, $P(CW|x)$ is also computed and saved, although that takes a small fraction of time compared to the big likelihood table used for inference.

Then before the generative agent can generate a response, a stimulus is needed. This stimulus comes from the random or adaptive stimulus selection which uses $P(\theta)$ and not yet $P(\theta|x_s, y_s)$, as prior. This stimulus is then given to the generative agent, which looks up the probability in its likelihood table and generates a response with it. This response is then given to the inference program as data alongside the stimulus that was given to the generative agent. This continues until one experiment is over. The posterior parameter combination distribution is then $P(\theta|x_t, y_t)$. This is used for the analyses and plots.

2.6.1 Model Parameter Recovery

To test the effect of adaptive stimulus selection, the young generative model’s parameters will be recovered using adaptive stimulus selection. The parameter values of the young generative model come from [1]. With a maximum bias of $\pm 5.8^\circ$ at a frame orientation of $\pm 22.3^\circ$ and the shape of the patient generative model, the bias seems to be a combination of the biases from both the old and patient generative models.

First, the individual parameters are recovered while setting the other parameters to their generative values. Then pairs of parameters will be recovered and the negative log likelihood will be used to illustrate the recovery. Finally, the model with all free parameters will be recovered. The experiment will be run 10 times with 500 iterations for each condition to get a more informative distribution of the recovered parameter values. The number of iterations needed to be

lower than in [1], which had 1620 iterations per condition, to convince ourselves that adaptive stimulus selection is a worthwhile endeavour. If the results show comparable performance with fewer than a third of the original iterations, it can be concluded that adaptive stimulus selection is worth pursuing.

	κ_{ver}	κ_{hor}	τ	κ_{oto}	λ
Min	32.53	0.4056	0.58	133.4	0.00
Max	176.7	77.17	1.0	156.5	0.06

Table 3: Ranges of the parameters used in recovery of the young generative model’s parameters. The minimum and maximum values are the mean parameter values in [1] minus, respectively plus, 2 standard deviations. 95% of the data now falls within these ranges.

The parameter value ranges can be seen in Table 3. The κ parameter ranges are discretized non-linearly. In [1], the mean and standard deviation σ values are given. From these, the minimum and maximum σ values are computed as the mean value minus, respectively plus, 2 standard deviations. This σ range is then discretized linearly into 10 or 25 steps. Then this linear σ range discretization is translated to a non-linear κ range discretization, by translating every individual σ value to a corresponding κ value. The κ range discretization is then non-linear because the relation between σ and κ is non-linear ($\kappa \approx 1/\sigma^2$). The ranges of τ and λ are discretized linearly.

When recovering a single parameter or a pair of parameters, the range is discretized into 25 steps. When recovering the parameters of the full model, the range is discretized into 10 steps. Increasing that number further would make the inference too expensive for the available hardware. The average trial time was still below 1 second, but running 20 experiments with a total of 10,000 iterations to generate the data for the figures, took 3 hours. The number of steps could easily be increased in the future with more available computation power.

The single parameter and parameter pair recoveries are a sanity check to see whether the model’s parameters are recoverable before trying to recover the model with all free parameters. The analyses for the single parameter and parameter pair inference were mainly used to make a better inference program, but they still give a good insight and are worthy of discussion.

3 Results

3.1 Adaptive Stimulus Selection

The selected stimulus orientations for both the rods and frames during the full model experiment can be seen in Figure 4.




Figure 4: Selected stimulus orientations using random and adaptive stimulus selection for 500 trials. These data were gathered running the model once with all free parameters.

The adaptive stimulus selection plot shows that the procedure selects rod and frame orientations with quite a big variability at the start of the experiment, but that the procedure is quite selective at the end of the experiment. In other words, the adaptive stimulus selection procedure converges to a small number of rod-frame pairs. These are $(7^\circ, -10^\circ)$ and $(7^\circ, 25^\circ)$. The $(7^\circ, 25^\circ)$ pair can be explained by the bias of the generative model. 7° is closest to the maximum of the bias and 25° is closest to the frame orientation at which that bias is experienced. So adaptive stimulus selection chooses the frame orientation that exhibits the largest bias and the rod orientation with which it is the easiest to ascertain the size of said bias.

Besides the size of the bias, the response variability also needs to be recovered in order to recover the parameters of the model. Selecting stimuli where the probability of responding clockwise or counterclockwise is more than 99% is not informative for inferring the model parameters. Probabilities of 50% are also not informative for learning the model parameters, so adaptive stimulus selection will try to select the stimuli where the probability is approximately 80%. That is the reason no frames that are oriented upright or close to upright are selected; the information is not informative for learning the response variability of the generative model.

3.2 Single Parameters

The parameter value distributions for single parameters can be seen in Figure 5. The experiment was run 10 times for both random and adaptive stimulus selection to get an accurate mean and standard deviation. Each experiment had 500 iterations.

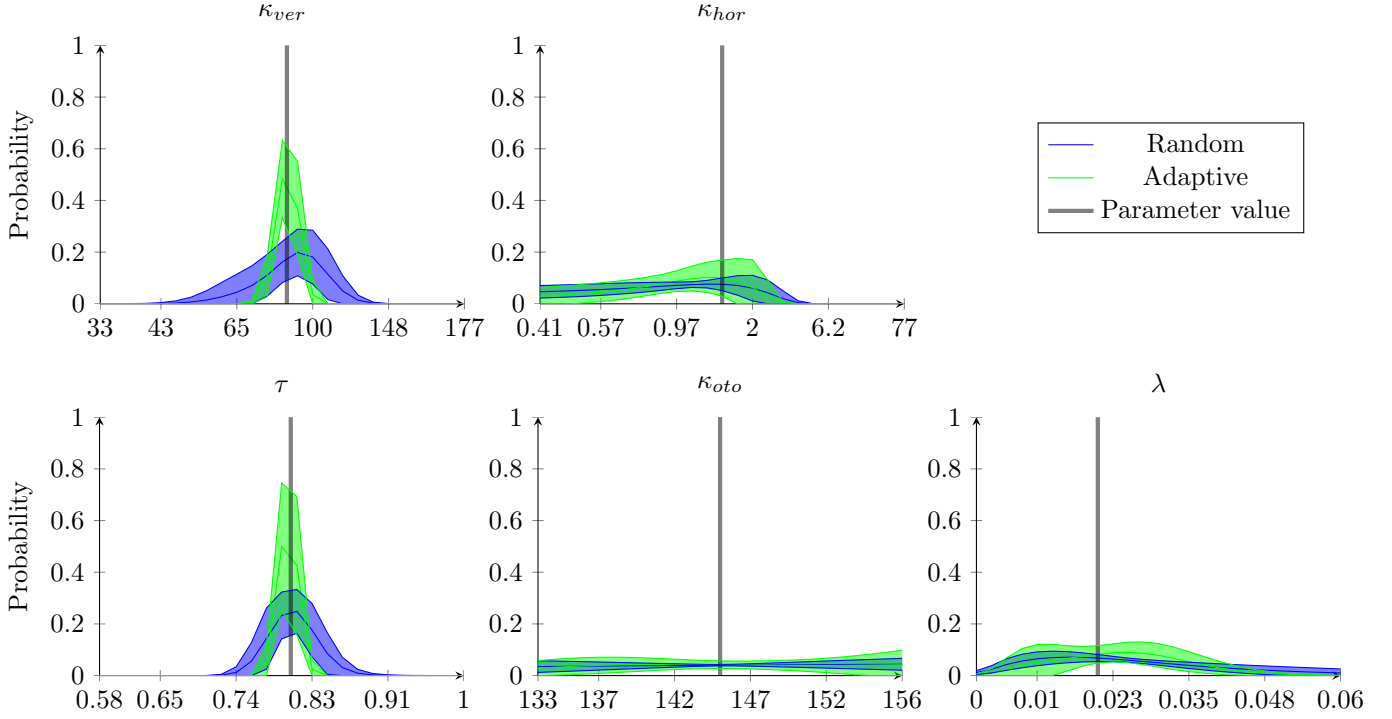


Figure 5: The single parameter value distributions at the end of the experiment. The surface is between the mean minus the standard deviation and the mean plus the standard deviation. All the parameter ranges are discretized into 25 steps. The generative parameter value can be seen by the vertical grey line.

A clear difference between random and adaptive stimulus selection can be seen for κ_{ver} and τ . The distribution from adaptive stimulus selection is more precise and more accurate than the distribution from random stimulus selection. This also holds for κ_{hor} to a lesser extent. The standard deviation of the probabilities of the adaptive distributions is higher compared to the random distributions. This makes sense, because the random experiments have roughly the same total stimuli set after each experiment, whereas the adaptive experiments have stimuli based on the responses, which is different for each experiment. So the standard deviation is higher, because there is more variety in the total set of stimuli for the adaptive compared to the random experiments.

A different story can be told about the distributions of κ_{oto} and λ . The distribution of κ_{oto} is basically flat and no significant MAP value can be picked for both the adaptive and random distributions. This can be explained by the small range of κ_{oto} . Compared to κ_{ver} and κ_{hor} , κ_{oto} 's range is much smaller, despite the range being between minus and plus 2 standard deviations from the generative parameter value. So any value it picks in this range gives satisfactory results and that is why the distribution is so flat.

The distribution for λ does not tend towards the generative parameter value and it is also relatively flat in comparison to κ_{ver} , κ_{hor} and τ . This can simply be an inherent consequence of the lapse parameter. The difference in responses between a lapse rate of 0.01 and 0.02 is minuscule and far more responses than the 500 in the experiment would be needed to get a more precise distribution. But even for the λ distribution it holds that the mean probability of the MAP value is higher for the adaptive compared to the random distribution and the standard deviation is also higher.




Figure 6: The normalized standard deviations of the marginalized parameter value distributions for each trial during the experiment. Each model had only one free parameter. The experiment was run 10 times for both random and adaptive stimulus selection to get an accurate mean and standard deviation. The surface is between the mean minus the standard deviation and the mean plus the standard deviation. All the parameter ranges are discretized into 25 steps. The surface is smaller than 0.30046, because that is the standard deviation of a normalized uniform distribution with 25 steps.

The standard deviations of the distributions in Figure 5 were normalized to be able to compare the precision of the distributions throughout the experiment across parameters. The results can be seen in Figure 6. For κ_{ver} , the normalized standard deviation is computed as follows. This generalizes to the other param-

eters as well. First, the parameter range discretization is normalized to a linear discretization from 0 to 1 ($\kappa_{ver,norm}$). Then the parameter value distribution of κ_{ver} is computed from $P(\theta|D_t)$, just like for Figure 5. Third, the normalized parameter mean value is computed as $\overline{\kappa_{ver,norm}} = P(\kappa_{ver}|D_t) \cdot \kappa_{ver,norm}$. Lastly, the parameter distribution standard deviation is computed with

$$\sigma_{\kappa_{ver,norm}} = \sqrt{\int (P(\kappa_{ver}|D_t) \cdot (\kappa_{ver,norm} - \overline{\kappa_{ver,norm}})^2) d\kappa_{ver}}$$

Looking at Figure 6, the advantage of adaptive stimulus selection becomes extremely apparent. The standard deviations of the adaptive parameter distributions are significantly lower for all parameters at the end of the experiment when compared to the random parameter distributions.

As expected, the standard deviations for τ and λ have a smaller initial value due to the beta parameter prior. The initial standard deviation for the κ parameters is the maximum standard deviation, because these parameter had a uniform prior distribution.

The three recoverable parameters had the lowest mean standard deviations at the end of the experiment. κ_{ver} has converged after about 200 trials when running adaptively, whereas it converges after 400 trials when running randomly. τ converges even faster. It converges after about 150 trials adaptively and after about 300 trials randomly. The other three parameters have not converged yet at the end of the experiment.

Interestingly, the three recoverable parameters show a non-linear relation between trial number and standard deviation, whereas the unrecoverable parameters show a linear relation. So the initial increase in precision in the first trials is missing from the unrecoverable parameters.

3.3 Parameter Pairs

To investigate the recovery of parameter pairs, the negative log likelihood will be used in a so called corner plot. A contour plot will be presented for each parameter pair, see Figures 7 and 8.

Random Stimulus Selection

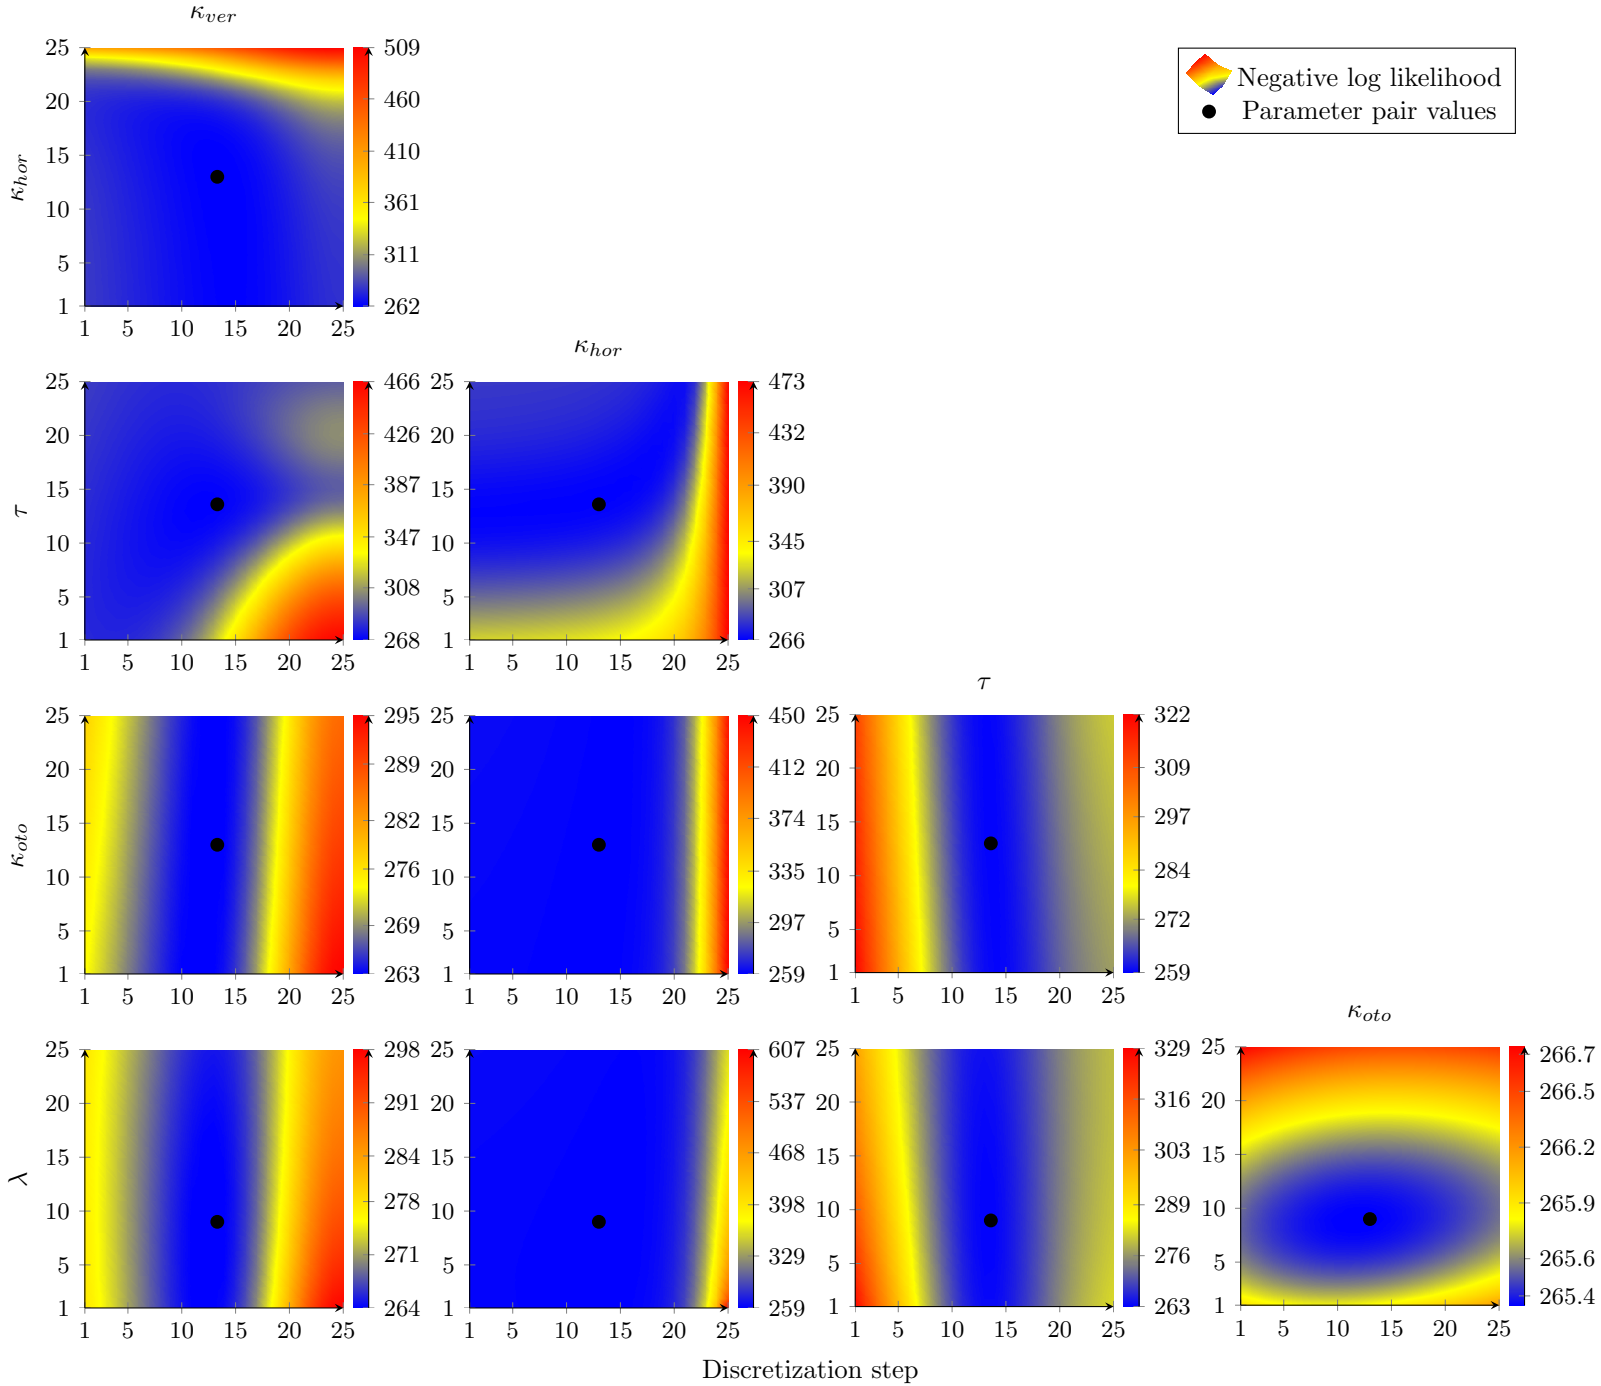


Figure 7: Mean negative log likelihood using random stimulus selection. The black marks stand for the generative parameter pair values. Each parameter range is discretized into 25 steps and the experiment was run 10 times with 500 iterations each time to get an accurate mean. The color bar next to a plot shows the negative log likelihood value associated with each color.

Adaptive Stimulus Selection

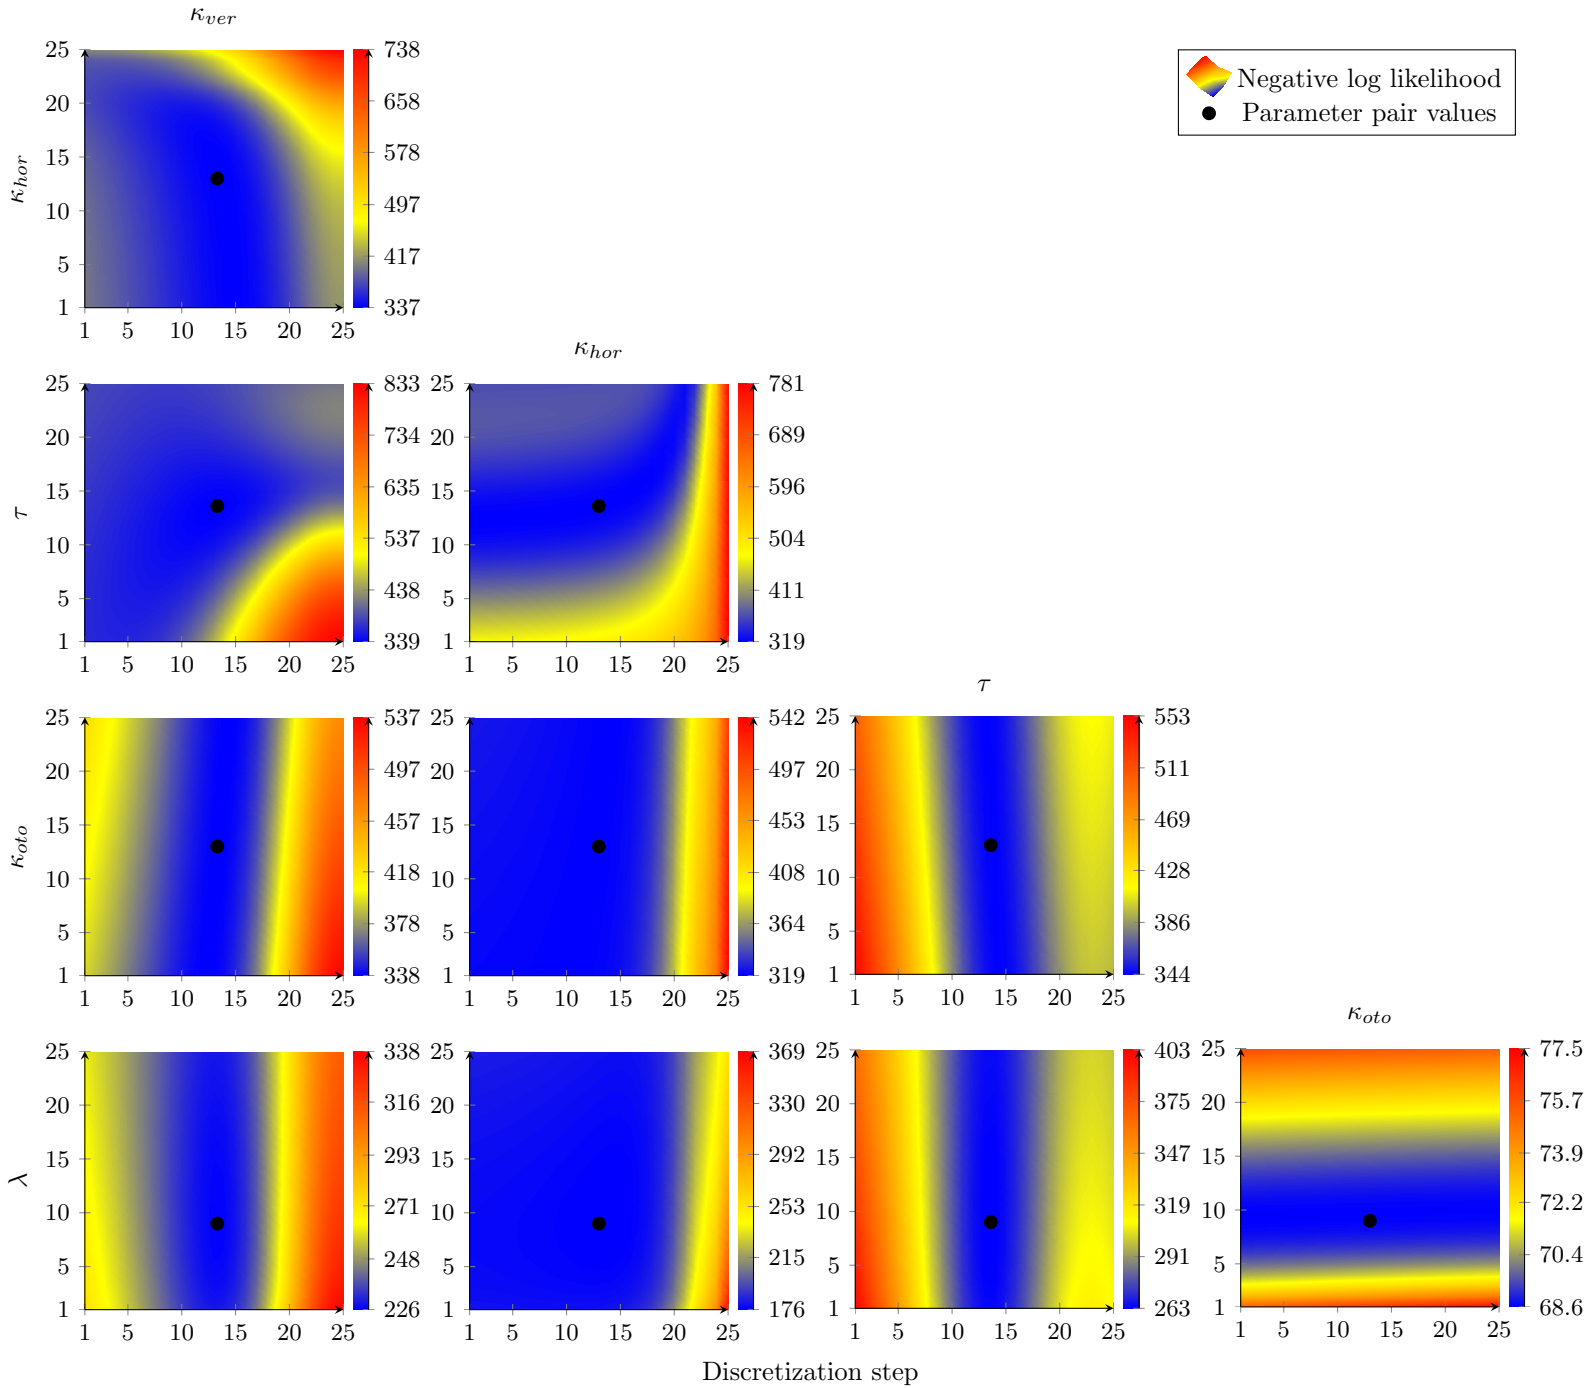


Figure 8: Mean negative log likelihood using adaptive stimulus selection. See Figure 7 for more information on the data.

What should be immediately apparent is that the negative log likelihood surfaces are almost identical between random and adaptive stimulus selection. In fact, these figures give little to no information about the benefit of adaptive

stimulus selection on recovering parameter pair values.

However, a lot can be said about how much the model’s parameters can be recovered. As already seen in Figure 5, κ_{oto} and λ were not recovered very well on their own. So when these parameters are paired with another parameter, the parameter pair negative log likelihood could also not recover these parameters very well. Take for example the negative log likelihood of κ_{ver} and κ_{oto} . Given a κ_{oto} value, the change in negative log likelihood when choosing different values for κ_{ver} is big. In contrast, when given a κ_{ver} value, the change in negative log likelihood when choosing different values for κ_{oto} is small. So only the chosen value for κ_{ver} seems to have an effect on the negative log likelihood and any κ_{oto} will do. So κ_{oto} could not be recovered in combination with κ_{ver} because it had no effect on the negative log likelihood of the parameter pair.

Do be cautious when making conclusions using Figures 7 and 8. The color maps are not normalized. When, for example, looking at the random negative log likelihood of the κ_{oto} and λ pair, it might be concluded that this parameter pair was successfully recovered; the negative log likelihood valley has an elliptical shape and both κ_{oto} and λ have a noticable effect on the negative log likelihood. However, the difference between the minimum and maximum negative log likelihood values is only 1.4 approximately, see the range next to the colorbar of that plot. So the visual difference does not actually exist when looking further into the data.

To give an insight in the difference adaptive stimulus selection makes, the meta data of the negative log likelihoods is shown in Tables 4, 5 and 6.

Random Stimulus Selection

Mean \pm SD	κ_{ver}	κ_{hor}	τ	κ_{oto}
κ_{hor}	89.8 \pm 10.6 1.44 \pm 0.884			
τ	78.5 \pm 12.3 0.799 \pm 0.0325	2.16 \pm 1.85 0.832 \pm 0.0576		
κ_{oto}	80.2 \pm 10.0 145 \pm 10.2	1.62 \pm 0.65 145 \pm 9.89	0.804 \pm 0.0232 143 \pm 9.3	
λ	79.6 \pm 13.2 0.0128 \pm 0.01	1.91 \pm 0.936 0.02 \pm 0.0181	0.795 \pm 0.0235 0.022 \pm 0.0151	143 \pm 10.7 0.019 \pm 0.0189

Table 4: The mean and standard deviation parameter values at the minimum of the negative log likelihood surface using random stimulus selection. The experiment was run 10 times to get an accurate mean, each parameter range was discretized into 25 steps and the experiment was run for 500 iterations. The top two numbers in a cell belong to the column parameter and the bottom two numbers in a cell belong to the row parameter.

Adaptive Stimulus Selection

Mean \pm SD	κ_{ver}	κ_{hor}	τ	κ_{oto}
κ_{hor}	91.9 \pm 8.79	1.13 \pm 0.581		
τ	84.4 \pm 8.88	1.71 \pm 1.67	0.795 \pm 0.0112	0.809 \pm 0.0387
κ_{oto}	87.2 \pm 8.5	1.25 \pm 0.360	0.808 \pm 0.0157	145 \pm 6.31
λ	86.2 \pm 4.79	1.47 \pm 0.685	0.801 \pm 0.014	147 \pm 7.23
	0.024 \pm 0.0177	0.0245 \pm 0.014	0.0258 \pm 0.017	0.0218 \pm 0.0083

Table 5: The mean and standard deviation parameter values at the minimum of the negative log likelihood surface using adaptive stimulus selection.

What becomes apparent is the fact that the mean minimum points of the negative log likelihood surfaces are close to the generative parameter values. But, as was already shown by the previous two figures, the difference in negative log likelihood between the minimum point and points around the minimum point is often small, leading to unrecoverability of the model’s parameters.

The standard deviations in Tables 4 and 5 do show a difference between random and adaptive stimulus selection. Most standard deviations of the adaptive experiment are smaller than the standard deviations of the random experiment and the ones who are not, are of κ_{oto} and λ . In other words, for the parameters that are recoverable (κ_{ver} , κ_{hor} and τ), the standard deviation is always smaller of the adaptive experiment compared to the random experiment. This suggests that the adaptive experiment has converged more to some parameter pair values than the random experiment, which in turn suggests that adaptive stimulus selection has formed a more precise posterior than random stimulus selection.

The standard deviations of κ_{oto} might seem small compared to the mean value of κ_{oto} , but one needs to remember the small range of κ_{oto} . Compared to their ranges, the standard deviations of κ_{oto} and λ are relatively much larger than the standard deviations of the recoverable parameters.

Random					Adaptive				
Min-Max	κ_{ver}	κ_{hor}	τ	κ_{oto}	Min-Max	κ_{ver}	κ_{hor}	τ	κ_{oto}
κ_{hor}	262-509				κ_{hor}	337-738			
τ	268-466	266-473			τ	339-833	319-781		
κ_{oto}	263-295	259-450	259-322		κ_{oto}	338-537	319-542	344-553	
λ	264-298	259-607	263-329	265.4-266.7	λ	226-338	176-369	263-403	68.6-77.5

Table 6: The mean minimum and maximum value of the negative log likelihood surfaces using random and adaptive stimulus selection, respectively.

To give more evidence toward the positive effect of adaptive stimulus selection on the recoverable parameters, the minimum and maximum negative log likelihood values are shown in Table 6.

For all parameter pairs except κ_{hor} - λ , the difference between the minimum and maximum negative log likelihood values is bigger for the adaptive experiment compared to the random experiment. In other words, the parameter pair values around the minimum values have a relatively higher negative log likelihood for the adaptive compared to the random experiment. So again, the negative log likelihood is in general more precise for the adaptive experiments.

Another curiosity is that the minimum value for all parameter pairs where λ is not included is smaller in the adaptive experiment. This is in contrast to the parameter pairs where λ is included, where the minimum value is lower for the adaptive compared to the random experiment. This might be explained by the fact that different stimuli were selected to create these surfaces in the random and adaptive experiments. Because of that, the minimum values of the surfaces cannot be compared and the curiosity just accidentally came out of the experiments.

3.4 Causes for Unrecoverability

A number of settings were changed to try to understand why the κ_{oto} and λ parameters were unrecoverable. The first that was tried was more iterations. The result can be seen in Figure 9.




Figure 9: The mean final parameter value distributions of an experiment with 5000 iterations. See Figure 5 for more details on the data.

With 10 times more iterations, the recovery of κ_{oto} is a lot better compared to Figure 5, yet still unsatisfactory. There is a clear difference between adaptive and random stimulus selection, but the distribution is still unprecise compared to the distributions of κ_{ver} and τ in Figure 5.

λ was recovered a lot better with a larger number of iterations. The MAP value of the distribution is close to the generative parameter value and the distribution is precise enough to have a certain degree of confidence on its MAP value. There is also a clear difference between adaptive and random stimulus selection, but this was already the case with 500 iterations.

With a lapse rate of 0.02, the subject will lapse, on average, once every 50 iterations. This means that, in the first experiment with 500 iterations, the subject only lapsed an average of 10 times. With 5000 iterations, the subject will lapse, an average of 100 times. So it makes sense that λ is recoverable with more data; an average of 10 iterations where the subject lapsed, was simply not enough to learn the value of λ .

It must be noted that 5000 iterations is approximately 3 times the number of iterations in [1]. So the distributions in Figure 9 will never be achieved in an empirical experiment.




Figure 10: The mean final parameter value distributions of an experiment where the parameter ranges are discretized into 10 steps instead of 25. See Figure 5 for more details on the data.

When using fewer discretization steps, the probabilities are artificially increased, because the probability mass is now divided over fewer data points. But if it could not recover the parameter before, it cannot recover the parameter now. Note though, λ has quite an informative prior and the data presented in Figure 10 still closely follows the shape of that prior. The data in Figure 10 is basically the data in Figure 5 multiplied by $25/10 = 2.5$.




Figure 11: The mean final parameter value distributions of an experiment with ranges between minus and plus 20 standard deviations from the generative value. Just like with the original ranges, negative values were excluded from the ranges. See Figure 5 for more details on the data.

Increasing the range does make the parameters recoverable, but at the cost of specificity. Figure 11 also gives more evidence that any value within the

original range of κ_{oto} will suffice and give satisfactory results. The width of the distribution of κ_{oto} is still surprisingly big, considering the range has been increased 10-fold. The advantage of adaptive stimulus selection has now become apparent for κ_{oto} with the same properties as previously. The range of λ is skewed because of the minimum value of 0. The new range goes from 0 to 0.42 and no human has such a high lapse rate. So it is not surprising that λ could be recovered with this increased range. But the results are not usable, because the resolution of the range is so low.

The original ranges cover 3.3 and 3.4 discretization steps for κ_{oto} , respectively λ . So this scenario not only increases the range, but also decreases the number of interesting parameter values, generating the same effect as in Figure 10. The probabilities are artificially higher and still closely follow the shape of the prior. The only difference in the latest plot is the decreased performance of random stimulus selection compared to adaptive stimulus selection for κ_{oto} . The difference between adaptive and random stimulus selection is bigger with a wider range, but that is simply because random stimulus selection is underperforming compared to the original range of κ_{oto} .

In conclusion, how can we make κ_{oto} and λ recoverable? More data! But the whole purpose of this project is to decrease the number of data points necessary to reliably recover the vestibular model parameters. So we have arrived at an impasse; we cannot increase the number of iterations, because that would defeat the purpose of the project, but then we cannot recover two parameters of the model, a purpose of the project as well. In the project, the priority is higher for adaptive stimulus selection and a low number of iterations, so the number of iterations will be kept at 500, accepting the unrecoverability of κ_{oto} and λ .

3.5 All Free Parameters

Next, the model with all free parameters was tested. See Figure 12 for the final parameter value distributions and Figure 13 for the standard deviations of these distributions throughout the experiment.

It has become extremely clear that the model’s parameters are not recoverable when all the parameters are free. All the MAP values of the parameters, except for τ , are far from the generative parameter values and the distributions for the κ parameters are relatively flat compared to the other two distributions. Furthermore, the difference between random and adaptive stimulus selection is negligible.

This was to be expected, given the unrecoverability of κ_{oto} and λ in Figure 5 and the negative log likelihoods of Figures 7 and 8 with their flatness around the minimum of the negative log likelihood. But it is still disappointing that the full model’s parameters are unrecoverable as well.

While trying to increase the performance of the model parameter inference, the priors, spline interpolation and non-linear κ range discretizations were included in the inference of the model parameters. These additions did improve the inference performance, but it was not enough to get satisfactory performance. So it seems that the model’s parameters are inherently unrecoverable.

But the results are nonetheless interesting. Looking at Figure 12, the distributions for τ and κ_{oto} have unusual shapes. Furthermore, the parameters that were unrecoverable in Figure 5 are also unrecoverable in the full model, but their standard deviations would suggest otherwise. κ_{ver} and κ_{hor} were recoverable when recovering single parameters, but in the full model, they are not recoverable anymore. Clearly something interesting and peculiar is happening during the experiment to arrive at these final distributions.

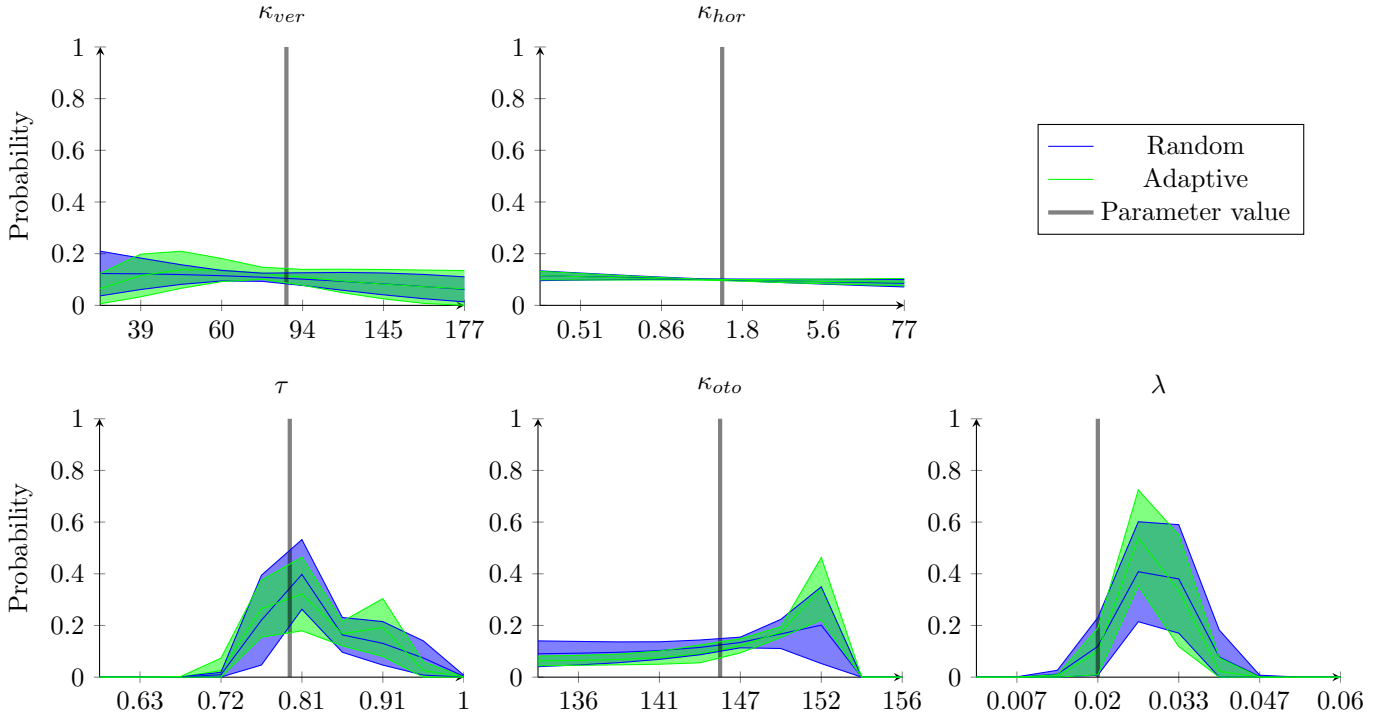


Figure 12: The marginalized parameter value distributions for the model with all free parameters at the end of the experiment. All the parameters ranges are discretized into 10 steps. See Figure 5 for more information on the data.

Looking at Figure 13, a few noteworthy observations will be noted. First of all, the initial standard deviations of τ and λ are lower than the κ standard deviations. This can be explained by the beta parameter priors for τ and λ , whereas the initial distributions for the κ parameters are uniform. Furthermore, the final standard deviations of τ and λ are also lower when compared to the κ standard deviations.

κ_{hor} did not change its standard deviation at all, this can also be seen in Figure 12, where the final distribution seems to be uniform for both random and adaptive stimulus selection.

The only parameter that shows a significant difference between random and

adaptive stimulus selection is λ . But because λ could not be recovered successfully, these data cannot be used to draw conclusions on the benefit of adaptive stimulus selection.

An interesting observation could be made from the standard deviation of κ_{ver} when running the adaptive experiment. During the first 200 trials, the standard deviation for the marginalized distribution of κ_{ver} is not different from a uniform distribution. Then suddenly from 200 trials, the mean standard deviation starts decreasing until it has decreased more than the mean standard deviation for κ_{ver} in the random experiment. This difference between the adaptive and random experiments could only be seen for κ_{ver} .

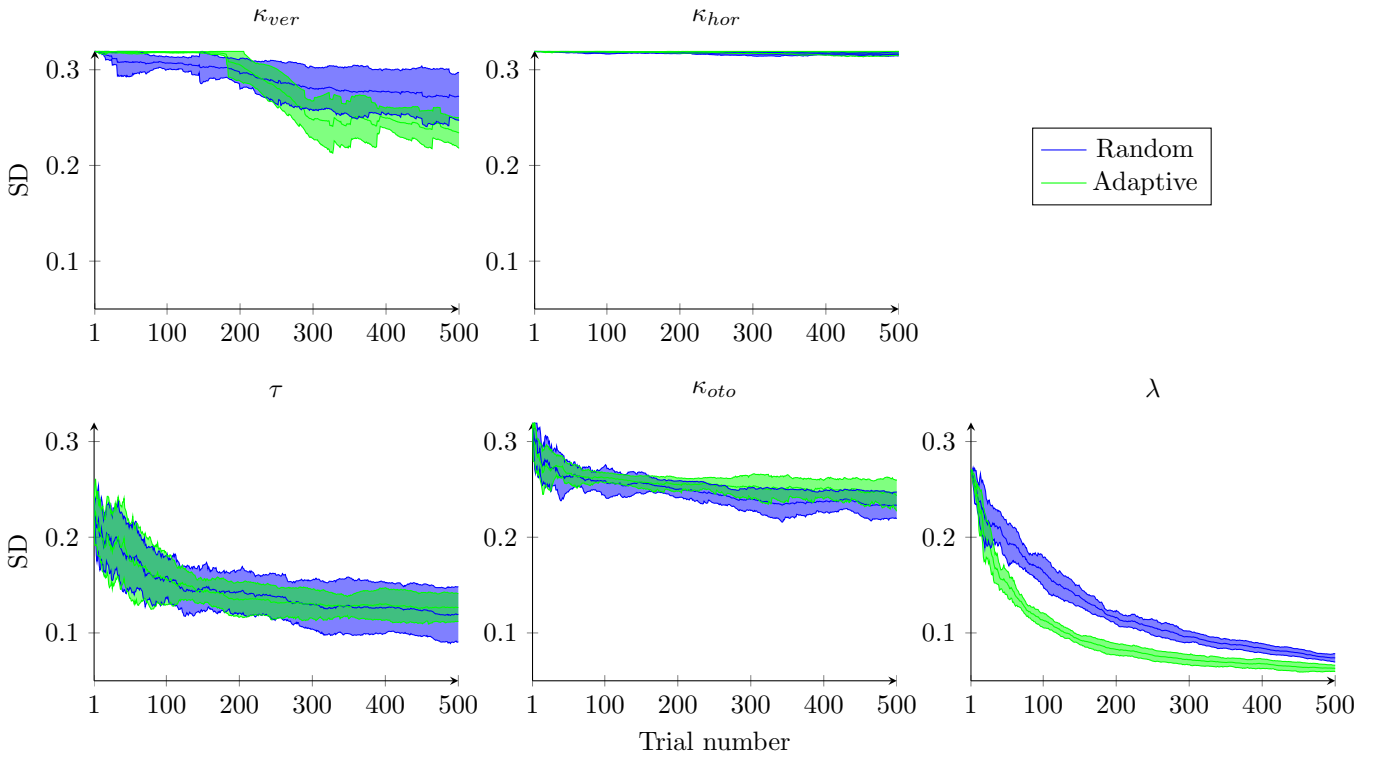


Figure 13: The normalized standard deviations of the marginalized parameter value distributions for each trial during the experiment. The model with all free parameters was used to make this figure and all the parameter ranges were discretized into 10 steps. See Figure 6 for more information on the data.

4 Discussion

The main goal of the present thesis was to successfully recover the proposed visual-vestibular model's parameters. On top of that, the advantage of adaptive

stimulus selection over random stimulus selection was to be shown during the model parameter recovery. Single parameters and parameter pairs could partly be recovered, but recovering the parameters of the full model was unsuccessful. However, the effect of adaptive stimulus selection compared to random stimulus selection has been extensively shown when single parameters and parameter pairs were recovered. The priors for τ and λ , spline interpolation and the non-linear κ parameter ranges did improve the performance of model parameter recovery, but it did not add a leap in performance needed to successfully recover the model parameters.

κ_{oto} could not be recovered on its own. As already mentioned, the small range of κ_{oto} could be a reason for this unrecoverability. Each step of the range gives satisfactory results, so the inference program selects all steps as equally likely. If comparable ranges were made for the old and patient generative models, then the ranges for κ_{oto} would be wider compared to the range of the young generative model. As seen in Figure 11, κ_{oto} could be recovered with a wider range. So choosing a more liberal parameter value range to be used in the experiments, could be a first step to successfully recover κ_{oto} .

λ could also not be recovered as a single parameter. As already stated, this can be explained by the minuscule change in responses when the λ value changes. Furthermore, with a lapse rate of, for example, 0.03, 3% of the responses will be caused by a lapse of the subject. So in an experiment with 500 iterations, 15 iterations are informative for learning the lapse rate of such a subject. The fact that very few iterations are informative for the lapse rate and the fact that a change in the lapse rate has almost no effect on the responses, gives a good explanation as to why λ , on its own and in combination with other parameters, could not be recovered.

Learning the parameter values of the full model was also unsuccessful. The vestibular likelihood was determined by the κ_{ver} , κ_{hor} and τ parameters. The τ parameter was recovered, whereas the κ_{ver} and κ_{hor} parameters were not successfully recovered. This can be explained by the way κ_{ver} , κ_{hor} and τ interact to make the vestibular likelihood. The τ parameter has a more direct effect on the vestibular likelihood, which means that changing the value for τ , changes the vestibular likelihood noticeably. This in turn changes the responses noticeably. κ_{ver} and κ_{hor} have an indirect effect on the vertical and horizontal vestibular likelihood precisions and so, changing these parameters has a smaller effect on the responses, compared to changing τ . This, in turn, means that recovering τ is easier than recovering κ_{ver} and κ_{hor} and that is why τ was recovered and κ_{ver} and κ_{hor} were not.

Fixing the head position of the subject might have been detrimental for the recovery of the model parameters. Including the head-in-space orientation in the stimulus, might have greatly helped the recovery of the model parameters. Adding an extra element to the stimulus gives adaptive stimulus selection an extra opportunity to select the most informative stimulus for learning the model parameters. So the parameters of the original model by [1] might actually be recoverable, but because the head orientation is fixed for the present thesis, the proposed model's parameters are not recoverable.

In Figures 12 and 13, the λ parameter has converged at the end of the experiment, whereas none of the other parameters have converged. Clearly, a lot of effort by the adaptive stimulus selection procedure is put into recovering λ . Perhaps if λ is fixed to 0.0, i.e. excluding the lapse rate from the model, or to the generative value, more effort is put into recovering the other four parameters. The lapse rate has arguably the smallest effect on the responses, so removing λ from the model does not fundamentally change it. Doing that might even make the other parameters better recoverable. τ could also be fixed at the generative parameter value, to see whether κ_{ver} and κ_{hor} would be better recoverable if the value for τ cannot change. A last option would be to propose a new visual-vestibular interactions model, whose parameters are recoverable. However, that would be outside the scope of the present thesis.

In the current paper, adaptive stimulus selection was used based on the minimization of the expected entropy. Other adaptive stimulus selection techniques also exist. One can, for example, select the stimulus that will be the most informative given all the previous trials and the potential next n trials [5]. That way, instead of selecting the stimulus that is best for the next trial, the stimulus is selected that is best for the next n trials. This technique could potentially speed up the model parameter recovery even more, to the point of clinical applicability.

It was assumed that the non-linear κ ranges behaved the same as the linear σ ranges in [1] and [9]. However, this assumption does not completely hold. It is most noticeable for κ_{hor} , because the range of that parameter is so wide. Look, for example, at the distribution of κ_{hor} in Figure 5. There is essentially no probability mass from step 19 and onwards. When it is expected that the κ range behaves like a linear σ range, then that distribution should be symmetric. So more investigation could be done into the effect of the range of the κ parameters.

No empirical experiments were performed for the present thesis. The generative model has been completely in silico. The proposed average trial time of 1 second for the model with all free parameters, might be too slow for empirical experiments. The code is written in Python, with extensive use of the NumPy library. There do exist NumPy alternatives, such as CuPy that support the same operations NumPy provides on the GPU. That greatly decreases the processing time needed. If such an alternative was used, the number of discretization steps in the model with all free parameters could be increased, while maintaining the same trial time. The choice could also be made to reduce the trial time and maintain the number of steps for the model with all free parameters.

The old and patient generative models have been presented in the Methods section, but they were left unused for the remainder of the thesis. The original goal of the thesis was to dissociate the old from the patient generative model given a series of stimulus-response pairs. These generative models are still in the thesis as a proper visual representation of vestibular functioning in young, old and patient subjects. But because the parameters of the young generative model could not be recovered, the original goal was abandoned to try to recover the young generative model's parameters.

Because of the unrecoverability of the model's parameters, Figures 7 and 8

are not very usable. The scale of the negative log likelihood surface is different for each plot in the figures and the differences at the minimum of the negative log likelihood are hard to discern. To remedy this problem, the color bars and tables were added. This does give the data that is needed, but it is not a nice visual representation anymore. Hence the call for caution on making conclusions using these figures.

5 Conclusion

The visual-vestibular interactions model's parameters could not be recovered. However, the effect of adaptive stimulus selection has been effectively shown. A number of visual representations of the data have been presented, to hopefully give the impression that the unrecoverability of the model's parameters is fundamental. In other words, the unrecoverability is due to the inherent properties of the model and not the presented inference technique.

In order to some day arrive at a clinically applicable vestibular functioning test, a different model must be proposed, whose parameters are recoverable. The techniques posed in the present thesis might be a guideline on the parameter recovery of that newly proposed model.

References

- [1] Alberts B. B. G. T.; de Brouwer A. J.; Selen L. P. J.; Medendorp W. P. "A Bayesian Account of Visual–Vestibular Interactions in the Rod-and-Frame Task". In: *eNeuro* 3.5 (2016), pp. 1–14. DOI: [10.1523/eneuro.0093-16.2016](https://doi.org/10.1523/eneuro.0093-16.2016). URL: <http://www.eneuro.org/content/3/5/ENEURO.0093-16.2016>.
- [2] Verhagen W.; Bom S.; Huygen P.; Fransen E.; Van Camp G.; Cremers C. "Familial Progressive Vestibulocochlear Dysfunction Caused by a COCH Mutation (DFNA9)". In: *Archives of Neurology* 57.7 (2000), pp. 1045–1047. DOI: [10.1001/archneur.57.7.1045](https://doi.org/10.1001/archneur.57.7.1045). URL: <https://jamanetwork.com/journals/jamaneurology/fullarticle/776853>.
- [3] Graybiel A.; Kerr W. A.; Bartley S. H. "Stimulus Thresholds of the Semi-circular Canals as a Function of Angular Acceleration". In: *The American Journal of Psychology* 61.1 (1948), pp. 21–36. DOI: [10.2307/1417289](https://doi.org/10.2307/1417289). URL: https://www.jstor.org/stable/1417289?seq=16#metadata_info_tab_contents.
- [4] Bagust J. "Assessment of Verticality Perception by a Rod-and-Frame Test: Preliminary Observations on the Use of a Computer Monitor and Video Eye Glasses". In: *Archives of Physical Medicine and Rehabilitation* 86.5 (2005), pp. 1062–1064. DOI: [10.1016/j.apmr.2004.05.022](https://doi.org/10.1016/j.apmr.2004.05.022). URL: <https://www.ncbi.nlm.nih.gov/pubmed/15895360>.

- [5] Kim W; Pitt M; Lu Z.; Myung J. “Planning Beyond the Next Trial in Adaptive Experiments: A Dynamic Programming Approach”. In: *Cognitive Science* 41.8 (2016), pp. 2234–2252. DOI: [10.1111/cogs.12467](https://doi.org/10.1111/cogs.12467). URL: <https://onlinelibrary.wiley.com/doi/full/10.1111/cogs.12467>.
- [6] Angelaki D.; Cullen K. “Vestibular System: The Many Facets of a Multimodal Sense”. In: *Annual Review of Neuroscience* 31.1 (2008), pp. 125–150. DOI: [10.1146/annurev.neuro.31.060407.125555](https://doi.org/10.1146/annurev.neuro.31.060407.125555). URL: <https://www.annualreviews.org/doi/abs/10.1146/annurev.neuro.31.060407.125555>.
- [7] Zupan L. H.; Peterka R. J.; Merfeld D. M. “Neural Processing of Gravito-Inertial Cues in Humans. I. Influence of the Semicircular Canals Following Post-Rotatory Tilt”. In: *Journal of Neurophysiology* 84.4 (2000), pp. 2001–2015. DOI: [10.1152/jn.2000.84.4.2001](https://doi.org/10.1152/jn.2000.84.4.2001). URL: <https://www.ncbi.nlm.nih.gov/pubmed/11024093>.
- [8] Alberts B. B. G. T.; Selen L. P. J. Medendorp W. P. “Age-related reweighting of visual and vestibular cues for vertical perception”. In: *Journal of Neurophysiology* 121.4 (2019), pp. 1279–1288. DOI: [10.1152/jn.00481.2018](https://doi.org/10.1152/jn.00481.2018). URL: <https://physiology.org/doi/abs/10.1152/jn.00481.2018?af=R>.
- [9] Alberts B. B. G. T.; Selen L. P. J.; Verhagen W.; Pennings R; Medendorp W. P. “Bayesian quantification of sensory reweighting in a familial bilateral vestibular disorder (DFNA9)”. In: *Journal of Neurophysiology* 119.3 (2018), pp. 1209–1221. DOI: [10.1152/jn.00082.2017](https://doi.org/10.1152/jn.00082.2017). URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5899303/>.
- [10] Girshick A. R.; Landy M. S.; Simoncelli E. P. “Cardinal rules: visual orientation perception reflects knowledge of environmental statistics”. In: *Nature Neuroscience* 14.7 (2011), pp. 926–932. DOI: [10.1038/nn.2831](https://doi.org/10.1038/nn.2831). URL: <https://www.nature.com/articles/nn.2831>.
- [11] Bak J. H.; Pillow J. W. “Adaptive stimulus selection for multi-alternative psychometric functions with lapses”. In: *Journal of Vision* 18.12 (2018), pp. 1–25. DOI: [10.1167/18.12.4](https://doi.org/10.1167/18.12.4). URL: <https://jov.arvojournals.org/article.aspx?articleid=2713945>.

Appendix

The code can be found at [GitHub](#). All the plots in the thesis can be easily made for one experiment at a time and the summary data across experiments can be printed to the terminal to replicate the presented data.