
© COPYRIGHTED BY

Manisha Bhardwaj

December 2013

**VISUAL DECISION MAKING IN THE PRESENCE
OF STIMULUS AND MEASUREMENT
CORRELATIONS**

A Dissertation

Presented to

the Faculty of the Department of Mathematics

University of Houston

In Partial Fulfillment

of the Requirements for the Degree

Doctor of Philosophy

By

Manisha Bhardwaj

December 2013

**VISUAL DECISION MAKING IN THE PRESENCE
OF STIMULUS AND MEASUREMENT
CORRELATIONS**

Manisha Bhardwaj

APPROVED:

Dr. Krešimir Josić (Committee Chair)
Department of Mathematics, University of Houston

Dr. Matthew Nicol
Department of Mathematics, University of Houston

Dr. Ilya Timofeyev
Department of Mathematics, University of Houston

Dr. Wei Ji Ma
Center for Neural Science and
Department of Psychology, New York University

Dean, College of Natural Sciences and Mathematics

Acknowledgements

I owe my deepest gratitude and intellectual debt to my advisor, Dr. Krešimir Josić, for his constant thrust of guidance, support, and patience. This dissertation would have been impossible without his enthusiastic encouragement and persistent help. His insightful comments and unsurpassed knowledge contributed enormously in shaping my scientific and mathematical thoughts. Frequent and regular hours of discussions with him have deeply influenced, and nurtured my academic and personal career. I am particularly grateful to him for his patient assistance in developing my technical writing and presentation skills. The financial support provided by him, and graciously offered several car rides are greatly acknowledged.

I extend my profound thanks to Dr. Wei Ji Ma for providing his insightful contributions and critical suggestions on my research projects. His useful critiques, meticulous comments, and area expertise greatly benefited the quality of my research work. His upcoming book and theoretical neuroscience course notes have served as excellent references for the introduction chapter in the dissertation. I am also thankful to him for offering the laboratory resources in running the experiments that are part of the dissertation.

My committee members Dr. Matthew Nicol and Dr. Ilya Timofeyev deserve a special note of thanks for generously serving on my defense committee, proof-reading this dissertation, and providing useful suggestions for its improvement.

I sincerely appreciate the contribution of my academic collaborators and colleagues for their valuable support and helpful suggestions on my research work. I am especially indebted to Dr. Ronald van den Berg for many discussions, numerous suggestions, and contributions in Chapters 4 and 5. I am also grateful to Sam Carroll for his support and valuable contributions in Chapters 6 and 7. I would also like to thank José Manuel López for technical support and programming help during my graduate years. I am also thankful to Dr. Claudia Pedroza of the University of Texas Health Science Center at Houston, for her valuable help and suggestions on Monte Carlo Markov chain methods.

I am grateful to the academic, financial, and administrative contribution of Department of Mathematics at the University of Houston and its staff in my development as a scholar and teacher. I express my greatest appreciation to Dr. David Blecher and Dr. Robert Azencott for writing me the helpful recommendations. I honor the invaluable contributions of all the professors and teachers who instructed me in various fundamental courses during my high school, undergraduate and graduate studies. I am specially thankful to the faculty members of the Mathematical Sciences Foundation at Delhi for motivating and preparing me to pursue a research career in mathematics.

Finally, I acknowledge the invaluable network of supportive, forgiving, generous, and loving friends without whom I could not have survived the process. All my friends in Delhi and Houston deserve a special note of thanks for their outstanding care and support. I am thankful to Sanat for helping me think rationally, questioning my ideas, and his eagerness to help me. I would like to thank my

childhood friend Suryansh for his brotherly love and affection. I am extremely fortunate to have Pankaj beside me with his unconditional support, care, and patience. I am also grateful to him for reviewing many chapters of the dissertation and mathematical discussions.

I dedicate my dissertation to my parents and family for their unconditional love, constant support, motivation, and belief in my abilities. My parents, Satvir Singh Bhardwaj and Pushpa Rani Bhardwaj, have always been the inspiration and driving force of my life. I am also blessed with loving and caring brothers, Nitin and Sumit, who have always supported me in my pursuits. Many thanks to my cousins and close relatives for their unwavering support and love. Cheerful, energetic, and lively conversations with them have kept my life balanced. This work would have been impossible without the love and patience of my family.

To my parents Satvir Singh Bhardwaj and Pushpa Rani Bhardwaj

**VISUAL DECISION MAKING IN THE PRESENCE
OF STIMULUS AND MEASUREMENT
CORRELATIONS**

An Abstract of a Dissertation
Presented to
the Faculty of the Department of Mathematics
University of Houston

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

By
Manisha Bhardwaj
December 2013

Abstract

Our brains process sensory information to infer the state of the world. However, the input from our senses is noisy, which may lead to errors in perceptual judgements. A number of theoretical studies have modeled perception as a process of probabilistic inference that involves making decisions based on uncertain evidence. Bayesian optimality is a general principle of probabilistic inference that has been successfully used to build quantitative models of perception. In addition, several experimental studies show that human observers make best possible decisions, and hence exhibit close to Bayes-optimal behavior on various visual perceptual tasks such as visual search, sameness judgement, and change detection. However, the impact of structured stimuli on decision-making remains largely unexplored. Moreover, the sensory measurements can themselves be strongly correlated to produce a structured representation of the stimulus input. These measurement correlations can interact with the structure of the external input in many possible ways and should not be considered in isolation.

In this work, we focus on visual search task to examine how visual perception is affected by structured input. We analyze the responses of subjects on a target detection experiment where the stimulus orientations were generated with varying strength of correlations across different experimental sessions. We fit several models to the experimental data using maximum-likelihood parameter estimation. We use rigorous model selection to find that human observers take into account stimulus correlations in detecting a target. However, they behave suboptimally in inferring the correct stimulus correlations that were used in the experiment. We

find that perhaps observers treat the partial stimulus correlations identically and behave differently when the stimuli are perfectly correlated.

We also describe how the relation between measurement and stimulus correlations affects the performance of an ideal Bayesian observer in a family of target detection tasks. We find that the effect of measurement correlations depends on its interaction with stimulus correlations and other statistical structure parameters. Measurement correlations always improves the performance of the ideal observer on a detection task with multiple targets; whereas in the case of single target, the impact is significant only in the presence of strong external structure.

Contents

1	Introduction	1
1.1	Perception as Bayesian inference	4
1.1.1	Bayes' theorem	5
1.2	Elements of Bayes' theorem	6
1.3	Bayesian modeling of perception	9
1.3.1	Visual and auditory perception	10
1.3.2	Psychophysical studies	11
1.3.3	Step 1: the generative model	12
1.3.4	Step 2: the inference or perception process	16
1.3.5	Step 3: the observer's estimate distribution	19
1.4	Visual search	21
1.4.1	An example of a target detection task	22
1.4.2	Generalizations	32
1.5	Outline of the dissertation	34
2	Stimulus correlations in a target detection task	38
2.1	Generative model	41
2.2	Inference process	45

CONTENTS

2.2.1	The log-likelihood ratio	46
2.2.2	Interpretation of the decision variable	52
2.3	MAP estimate distribution	52
2.4	Suboptimal models	53
2.4.1	Heterogeneous model, $\rho_s = 0$	54
2.4.2	Homogeneous model, $\rho_s = 1$	54
2.5	Experimental Methods	58
2.5.1	Subjects	58
2.5.2	Apparatus and stimuli	59
2.5.3	Sessions and blocks	60
2.5.4	Procedure	60
3	Model fitting and model comparison	63
3.1	Experimental data	64
3.2	Model predictions	66
3.2.1	Computing response probabilities	67
3.2.2	Monte Carlo algorithm	67
3.3	Maximum-likelihood estimation	69
3.4	Optimization techniques	71
3.4.1	Exhaustive or grid search	71
3.4.2	Genetic algorithm and pattern search	73
3.5	Psychometric curves	75
3.5.1	Types of psychometric curves	76
3.5.2	Predictions using synthetic data	78
3.5.3	Error measures	79
3.6	Model comparison	80
3.6.1	Bayesian model comparison (BMC)	81

CONTENTS

3.6.2	Bayesian information criterion (BIC)	83
3.6.3	Akaike information criterion (AIC)	84
3.7	Discussion	85
4	Data analysis I: model fitting	86
4.1	Psychometric curves of subjects data	88
4.1.1	Hit and false-alarm rates	89
4.1.2	Minimum target-distractor orientation difference	91
4.1.3	Sample standard deviation of distractor orientations	92
4.2	Models	93
4.2.1	Assumptions about ρ_s	95
4.2.2	Encoding precision assumptions	97
4.2.3	Summary of models	99
4.3	Equal precision models	102
4.3.1	Condition-independent precision J	102
4.3.2	Condition-dependent precision J	112
4.4	Variable precision models	120
4.4.1	Condition-independent mean precision \bar{J}	120
4.4.2	Condition-dependent mean precision \bar{J}	128
4.5	Need for model comparison	135
5	Data analysis II: model comparison	137
5.1	Model families	138
5.2	Equal versus variable precision	141
5.3	Condition-independent \bar{J} versus condition- dependent \bar{J}	144
5.4	Comparison based on $\rho_{s_{\text{assumed}}}$	148
5.5	Rejection rate analysis	152

CONTENTS

5.6	Parameter estimates	156
5.7	Conclusions and discussion	160
6	Measurement correlations in a single target detection task	163
6.1	Model description	166
6.2	Optimal observer theory	168
6.2.1	The log-likelihood ratio	171
6.2.2	Bayesian decision variable	175
6.2.3	Interpretation of the decision variable	176
6.2.4	Uncorrelated measurements, $\rho_x = 0$	178
6.3	Analysis and Results	179
6.3.1	Weak measurement noise, $\sigma_x^2 \ll \sigma_s^2$	180
6.3.2	Strong measurement noise, $\sigma_x^2 = \sigma_s^2$	187
6.4	Summary	190
7	Measurement correlations in a multiple target detection task	191
7.1	Model description	194
7.2	Optimal observer theory	196
7.2.1	Interpretation of the decision variable	201
7.3	Analysis and results	203
7.3.1	Weak external structure, $\rho_s < 1$	204
7.3.2	Strong external structure, $\rho_s = 1$	210
7.4	Mean stimulus orientation discrimination task	212
7.4.1	Inference process	214
7.4.2	Results	215
8	Discussion	219

CONTENTS

A	Notation table	227
B	Some mathematical results	229
B.1	Product and integral of normal distributions	229
B.1.1	Univariate normal distributions	230
B.1.2	Multivariate normal distributions	231
B.2	Determinant and inverse of a rank-1 matrix	232
B.3	Determinant and inverse of a rank- k matrix	232
C	Gabor Filter	234
	Bibliography	236

List of Figures

1.1	Bayesian modeling of a perceptual task - defining the generative model and deriving the inference process. (A) The generative model. The first step in Bayesian modeling is to define the generative model. This figure outlines the graphical representation of the generative model we will be using throughout the dissertation. The nodes represent the variables involved in the task and arrows determine the influence of one node on another. This influence is mathematically described in terms of conditional probabilities. The observer infers the (hidden) state of the world, W from the stimulus, s presented in the task by making a measurement, x of the stimulus. (B) Inference process. The second step in Bayesian modeling is to derive the inference process of an observer. That is, to understand the mathematical process by which the observer infers W based on the measurement, x . This step involves inverting the generative model and marginalizing over intermediate variable, s to compute a decision criterion and making an estimate, \hat{W} about the state of the world.	14
-----	---	----

- 1.2 **The process of marginalization in Bayesian modeling.** The generative model usually contains auxiliary variables that are not of primary interest, but they may have necessary information about the state of the world, W . Here the stimulus, s is an intermediate variable, but it links the world state variable of interest with the measurement, x . Marginalization is a process to deal with such ancillary variables to obtain the desired expression for the likelihood or posterior probability of the world state variable of interest. It involves averaging or integrating over the possible values of the ancillary variable and is very common in Bayesian modeling. . . . 19
- 1.3 **Steps involved in Bayesian modeling of a perceptual task.** The figure presents the schematic of a Bayesian inference process to model a perceptual task. We will follow this plan for all the tasks discussed in the dissertation. The first step of specifying the generative model involves describing the probability distributions to understand how sensory data are generated from the state of the world. The observer makes an estimate of the world state based on the sensory measurement on each trial of the task. This constitutes the second step of deriving the inference process in a Bayesian model. The estimate of the observer varies across trials in response to a fixed stimulus and follows a distribution. In the final step of Bayesian modeling of the task, this estimate distribution is computed. 20

- 1.4 **Bayesian modeling of a simple target detection task with two stimuli. (A) The generative model.** The binary variable, T describes the target presence in a trial. The two stimuli, s_1 and s_2 are chosen conditioned on T . When $T = 1$, one of the stimuli is a target with a vertical orientation, while the orientation of the other stimulus is chosen randomly from a normal distribution. The observer makes noisy and independent measurements, x_1 and x_2 of the two stimulus. **(B) The inference process.** The observer combines the two measurements to compute a decision variable, $d(x_1, x_2)$ and infers an estimate, \hat{T} of the world state variable, T . The decision variable, $d(x_1, x_2)$ is a log posterior ratio of the probability of reporting "target present" and "target absent", given the observer's measurements. If $d > 0$, the observer reports target is present and absent otherwise. **(C) Example displays in the task.** Since there are only two stimuli and one target, three types of visual displays can be presented to the observer. In the first two displays, the target is present to the left and right of the cross in the center. When there is no target, both stimuli are distractors and have randomly chosen orientations. The bottom display illustrates such an example. . . . 25
- 1.5 **Performance of an observer in a simple target detection example based on different models.** Proportion of correct responses as a function of the standard deviation of distractor orientations, σ_s for an observer based on different models. The optimal model has the maximum performance at all values of σ_s than other threshold models. A lower performance is observed at low standard deviations for all models since it becomes difficult to detect a target on the task among distractors that have relatively similar orientations to that of the target. As the standard deviation of the distractor orientations increases, the task becomes relatively easier and the performance increases for all models. The model with high threshold parameter, θ predicts a similar performance of the observer as that of the optimal model. This indicates that it is difficult to choose a model that is most consistent to describe the observer's behavior. . . . 31

- 2.1 **Statistical structure of relevant task variables in the optimal-observer model for a target detection task with stimulus correlations. (A) Generative model.** The nodes represent the variables in the task and arrows indicate conditional dependencies between them. The binary variable, T represents target presence for $T = 1$ and absent when $T = 0$. The standard deviation, σ_s and the pairwise correlation coefficient, ρ_s determine the structure of the stimulus, $\mathbf{s} = (s_1, s_2, \dots, s_N)$ in the task. An observer makes a measurement, x_i , of each presented stimulus s_i . These measurements are assumed to be noisy and independent between locations. **(B) Inference process.** The optimal observer infers T by "inverting" the generative model. The observer computes a decision variable, $d(\mathbf{x})$ based on the measurements, \mathbf{x} and it is given by the log-posterior ratio between the two possibilities, $\log(p(T = 1|\mathbf{x})/p(T = 0|\mathbf{x}))$. The sign of $d(\mathbf{x})$ gives the optimal estimate of T and it is denoted by \hat{T} 44
- 2.2 **Target detection experiment procedure. (A) Gabor patch.** Subjects were presented with 4 stimuli on each trial. Each stimulus was a Gabor patch. The figure illustrates an example of a Gabor patch from the experiment. **(B) Time course of a test trial.** The experiment started with a display of fixation cross in the center followed by the stimulus display for 0.1 sec. Subjects reported through a key press whether a vertical stimulus was present in the display. After their response, a display screen was shown with a green (correct) or red (incorrect) fixation cross to provide feedback. **(C) Time course of a practice trial.** Each experimental session started with 50 practice trials. The procedure of a practice trial was same as a test trial, except that an additional feedback screen was shown for 2 sec at the end of the trial. The extra display contained the original stimulus with a blue circle marking the target stimulus if it was present. **(D) Sample displays from different experimental conditions.** The experiment was divided in four different sessions. Each session was characterized by the unique value of the correlation coefficient $\rho_s \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$ used in generating stimuli. The order of the sessions was randomized across subjects. This figure shows example displays from each experimental session. 61

- 4.1 **Psychometric curves based on the experimental data.** Throughout the chapter, the error bars indicate unit standard error mean (s.e.m). **(A) Hit and false-alarm rates (left), and performance (right).** (Left) Hit (black) and false-alarm (red) rates as a function of correlation strength, ρ_s used in the experimental conditions. Hit (or false-alarm) rate shows a near-constant behavior in the first three experimental conditions and a large increase (or decrease) for $\rho_s = 1$. (Right) Mean subject performance in the four experimental conditions. **(B) Minimum target-distractor orientation difference.** Proportion of "target present" responses as a function of minimum difference between the target and any distractor, separately for target-present (left) and target-absent (right) trials in the four experimental conditions. Each curve corresponds to an experimental condition with $\rho_s \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$. Subjects' responses have a very similar trend for $\rho_s = 0, \frac{1}{3}$, and $\frac{2}{3}$, while their behavior is different when $\rho_s = 1$. **(C) Sample standard deviation between distractor orientations.** Proportion "target present" responses on different experimental conditions as a function of sample standard deviation between distractor orientations, in target-present (left) and target-absent (right) trials. The curves are similar to those in (B) except for $\rho_s = 1$ in target-absent trials (right). Since all distractors are identical in a trial when $\rho_s = 1$, there is only a single data point for the sample standard deviation. 90

- 4.2 **EP model 1 (J condition-independent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data.** Throughout the chapter, the shaded areas show the fit of the model and the circles are averaged subject responses. The error bars and shaded areas represent unit standard error of the mean for subjects' data and model fits, respectively. RMS errors and goodness of fit, R^2 values between the data and model predictions are indicated by blue and green numbers in the plots, respectively. **(A) Hit and false-alarm rates.** (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . **(B) Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). **(C) Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 104
- 4.3 **EP model 2 (J condition-independent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data.** **(A) Hit and false-alarm rates.** (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . **(B) Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). **(C) Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 105

4.4 **EP model 3 (J condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 107**

4.5 **EP model 4 (J condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 108**

4.6 **EP model 5 (J condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 110**

- 4.7 **EP model 6 (J condition-dependent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. (A) Hit and false-alarm rates.** (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . **(B) Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). **(C) Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 113
- 4.8 **EP model 7 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates.** (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . **(B) Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). **(C) Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 114
- 4.9 **EP model 8 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates.** (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . **(B) Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). **(C) Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 116

- 4.10 **EP model 9 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 118**
- 4.11 **EP model 10 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 119**
- 4.12 **VP model 1 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 122**

- 4.13 **VP model 2 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 123**
- 4.14 **VP model 3 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 125**
- 4.15 **VP model 4 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 126**

- 4.16 **VP model 5 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$ fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 127**
- 4.17 **VP model 6 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 129**
- 4.18 **VP model 7 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 131**

- 4.19 **VP model 8 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$ fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 132**
- 4.20 **VP model 9 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$ fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 133**
- 4.21 **VP model 10 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$ fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). 134**

5.1 **BIC and AIC model comparisons: equal versus variable precision models. (A) BIC model comparison.** (Left) Relative BIC differences of equal precision models with respect to variable precision models for individual subjects. Each vertical bar represents the BIC difference of an EP model from its equivalent VP model. (Right) Averaged (across subjects) relative BIC differences of equal precision models with respect to variable precision models. **(B) AIC model comparison.** Similar to (A) with AIC relative differences. A positive BIC or AIC relative difference indicates that the EP model is worse as compared to its equivalent VP model. Throughout the chapter, the error bars indicate the unit standard error mean (s.e.m). 143

5.2 **BIC and AIC model comparisons: VP models with condition-independent mean precision \bar{J} versus experimental condition-dependent \bar{J} . (A) BIC model comparison.** Relative BIC differences of condition-independent \bar{J} VP models with respect to condition-dependent \bar{J} models for each subject (left) and average subject (right). **(B) AIC model comparison.** Similar to (A) with AIC differences. 146

5.3 **BIC and AIC model comparisons: VP models with condition-dependent \bar{J} and different assumptions about $\rho_{s_{\text{assumed}}}$. (A) BIC model comparison.** Relative BIC differences of condition-dependent VP models having different assumptions about $\rho_{s_{\text{assumed}}}$ with respect to VP9 model (with $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$) for individual subjects (left) and average subject (right). **(B) AIC model comparison.** Relative AIC differences of condition-dependent VP models with respect to VP10 model (with $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$) for individual subjects (left) and average subject (right). 149

5.4 **Rejection rate curves: comparison of different model families in each model factor.** Each column corresponds to a factor and a color represents a particular model family or level belonging to the factor. The model comparison is based on all 20 models for each factor. **(A) Factor 1: precision.** Number of subjects for whom *all* models belonging to a certain family or level (EP or VP) are rejected as a function of the rejection criterion based on BIC (top) and AIC (bottom) differences. A model is rejected if it has a higher BIC or AIC than that of the winning model for a subject. For example, when both BIC and AIC rejection criteria is 100, all models of the EP family are rejected, while none of the subjects reject VP models. **(B) Factor 2: dependence of precision parameter on experimental conditions.** Similar to (A). **(C) Factor 3: assumption about correlation strength $\rho_{s_{\text{assumed}}}$.** Similar to (A). 154

5.5 **Individual subject parameter estimates of VP9 and VP10 models. (A) Parameter estimates of VP9 model.** Maximum-likelihood estimates of p_T , the observer’s prior probability of target presence (first column) and the scale parameter τ (second column) for each subject. Individual (gray) and average subject (black) estimates of mean precision \bar{J} in different experimental conditions (third column). Estimated values of correlation coefficient ρ_s for each individual (gray) and average subjects (black) as a function of true correlation strength (diagonal) in the experimental conditions. **(B) Parameter estimates of VP10 model.** Similar to (A). 159

- 6.1 **Statistical structure of relevant task variables in the optimal-observer model for a (single) target detection task with stimulus and measurement correlations. (A) Generative model.** The binary variable, T indicates target presence for $T = 1$, and absence when $T = 0$. The stimulus orientations, $\mathbf{s} = (s_1, s_2, \dots, s_N)$ are drawn from a multivariate normal distribution with mean vector, \mathbf{s}_D and covariance matrix, Σ_s . The standard deviation, σ_s and the correlation coefficient, ρ_s of distractor orientations determine the statistical structure of a visual scene. An observer makes measurements, $\mathbf{x} = (x_1, x_2, \dots, x_N)$ of the presented set of stimuli. These measurements are assumed to be unbiased, but correlated, and are drawn from a multivariate normal distribution with a covariance matrix, Σ_x . The correlation coefficient, ρ_x determines the extent of dependence between the sensory measurements. **(B) Inference process.** An ideal Bayesian observer computes the decision variable, $d(\mathbf{x})$ based on the measurements \mathbf{x} to make an estimate, \hat{T} of the true state variable, T . The decision variable, $d(\mathbf{x})$ is the log-posterior ratio between the two possibilities of making a response "target-present" or "target-absent", given the measurements and is given by $\log \frac{p(T=1|\mathbf{x})}{p(T=0|\mathbf{x})}$. The sign of $d(\mathbf{x})$ determines the estimate $\hat{T} = 1$ ($d(\mathbf{x}) > 0$) or $\hat{T} = 0$ ($d(\mathbf{x}) < 0$). 169

- 6.2 **Stimulus and measurement distributions in the single target detection with $N = 2$ stimuli for $\sigma_s = 15^\circ$ and $\sigma_x = 4^\circ$. (A) Stimulus distributions for $\rho_s = 0.5$.** Stimulus distributions on target present (left) and target absent (right) trials. On target present trials, one of the two stimuli is uniformly chosen to be the target and hence the distribution is aligned along the target orientation. On target absent trials, the two stimuli have a bivariate normal distribution. **(B) Measurement distributions for $\rho_s = 0.5$.** Measurement distributions for measurement correlations, $\rho_x = 0$ in response to the stimulus distributions in (A) on both target present (left) and target absent (right) trials. On target present trials, the measurement distribution is aligned towards the target orientation (left), while it represents the bivariate normal distribution in target absent cases (right). **(C) Overlap of measurement distributions for $\rho_s = 0.99$.** Measurement distributions on target present (orange) and absent (purple) trials are shown for $\rho_x = 0$ (top), and $\rho_x = 0.95$ (bottom). The overlap between the two measurement distributions, $\mathbf{x}|T = 1$ and $\mathbf{x}|T = 0$ reduces as the strength of measurement correlation, ρ_x increases and the two distributions become more distinguishable. Throughout the chapter, the axes are measured in terms of the standard deviation, σ , which is defined by $\sigma^2 = \sigma_s^2 + \sigma_x^2$ 170
- 6.3 **Performance of an ideal Bayesian observer in the presence of measurement and stimulus correlations on a single target detection task when the measurement noise is weak as compared to the external noise, $\sigma_x^2 \ll \sigma_s^2$ ($\sigma_s = 15^\circ, \sigma_x = 4^\circ$). (A) Variation in performance with measurement and stimulus correlations.** Performance of an optimal observer as a function of stimulus correlations, ρ_s and measurement correlations, ρ_x on a task with $N = 4$ stimuli. **(B) Change in performance with measurement correlations.** Proportion of correct responses as a function of measurement correlations, ρ_x in the case of weak external structure, $\rho_s = 0.5$ (left) and strong external structure, $\rho_s = 1$ (right) for $N = 4$ stimuli in the task. **(C) Decision boundary and measurement distributions for $N = 2$ stimuli and $\rho_s = 0.5$.** Decision boundary, $d_{\text{NST}}(\mathbf{x}) = 0$ (black) and measurement distributions corresponding to correct inferences (orange) and incorrect inferences (purple) on target present (left) and target absent (right) cases. The green curves (lines on left, ellipses on right) represent 2 units standard deviation of the stimulus distribution. 184

- 6.4 **Performance of an ideal Bayesian observer in the presence of measurement and stimulus correlations on a single target detection task when the measurement noise is comparable to the external noise, $\sigma_x^2 = \sigma_s^2$ ($\sigma_s = \sigma_x = 15^\circ$).** (A) **Variation in performance with measurement and stimulus correlations.** Performance of an optimal observer as a function of stimulus correlations, ρ_s and measurement correlations, ρ_x on a task with $N = 4$ stimuli. (B) **Change in performance with measurement correlations.** Proportion of correct responses as a function of measurement correlations, ρ_x for $\rho_s = 0$ (left) and $\rho_s = 0.5$ (right) on the task with $N = 4$ stimuli. (C) **Decision boundary and measurement distributions for $N = 2$ stimuli and $\rho_s = 0.5$.** Decision boundary, $d_{\text{NST}}(\mathbf{x}) = 0$ (black) and measurement distributions corresponding to correct inferences (orange) and incorrect inferences (purple) on target present (left) and target absent (right) cases. The green curves (lines on left, ellipses on right) represent 2 units standard deviation of the stimulus distribution. 189
- 7.1 **Performance of an ideal Bayesian observer in the presence of measurement and stimulus correlations on a multiple target detection task when the measurement noise is weak as compared to the external noise, $\sigma_x^2 \ll \sigma_s^2$ ($\sigma_s = 15^\circ, \sigma_x = 4^\circ$).** (A) **Variation in performance with measurement and stimulus correlations.** Performance of an optimal observer as a function of stimulus correlations, ρ_s and measurement correlations, ρ_x on a task with $N = 4$ stimuli and $n = 3$ targets. (B) **Change in performance with measurement correlations.** Proportion of correct responses as a function of measurement correlations, ρ_x for $\rho_s = 0$ (left), and $\rho_s = 0.5$ (right) on the task with $N = 4$ stimuli and $n = 3$ targets. (C) **Decision boundary and measurement distributions for $N = n = 2$ and $\rho_s = 0.5$.** Decision boundary, $d_{\text{NST}}(\mathbf{x}) = 0$ (black) and measurement distributions corresponding to correct inferences (orange) and incorrect inferences (purple) on target present (left) and target absent (right) cases. The green ellipses represent 2 units standard deviation of the stimulus distribution and the axes are measured in terms of the standard deviation, σ , which is defined by $\sigma^2 = \sigma_s^2 + \sigma_x^2$ 206

7.2 **Performance of an ideal Bayesian observer on the mean stimulus orientation discrimination task and multiple target detection task with $N = 4$ stimuli, $\sigma_s = 15^\circ$, $\sigma_x = 4^\circ$, and $\rho_s = 0.5$. (A) Mean stimulus orientation discrimination task.** Performance of an optimal Bayesian observer as a function of measurement correlations, ρ_x on the mean stimulus orientation discrimination task. **(B) Multiple target detection task.** Similar to (A) in a multiple target detection task with $n = 2$ 216

List of Tables

4.1	List of models fitted to the experimental data on the target detection task. Description of models fitted to the experimental data with different assumptions about encoding precision parameter and correlation coefficient. The number of free parameters per model is also listed.	101
5.1	List of models based on different factors and model families. Grouping of models in different model families based on a common assumption about the factor.	140
5.2	Maximum-likelihood parameter estimates of VP9 and VP10 models. The estimates of the mean precision and scale parameter of the precision distribution are given along with the estimates of correlation strength $\rho_{s_{\text{assumed}}}$ in each experimental condition (where applicable) and of p_T , the observer's prior probability that the target is present.	157
A.1	Mathematical Notations. Description of the mathematical notations used in the text.	228

Chapter 1

Introduction

Our decisions are based on sensory measurements and prior knowledge of the surroundings. The physical observations made by our eyes, ears, skin, and other sensory organs are transmitted to the brain. The brain integrates and interprets this information to draw inferences about the state of the world.

However, our sensory observations are typically incomplete and imperfect. They may not always reflect the true state of the world. Noisy and imprecise measurements can be difficult to interpret, and may lead to incorrect inferences. Thus, the question arises: how does the brain infer the state of the world from the inadequate, and uncertain sensory observations? Theoretical neuroscientists have hypothesized that our brain performs specific probabilistic computations to process the partially informative observations and makes an inference about the state of the world [69, 80, 160, 120, 90, 93, 100, 82]. Thus, perception can be thought of as a form of probabilistic inference [154, 65, 35]. The results of several experimental

studies indicate that the brain interprets sensory information probabilistically to make the best possible guess about the state of the world [43, 68, 10, 83, 84, 53, 54, 94, 151, 98, 75, 125]. These studies suggest that the brain computes the probability of many interpretations and chooses the one that is most likely. Such a strategy of framing the best possible perceptual inference is known as *optimal probabilistic inference* or *Bayesian inference*. The Bayesian theory of perception is based on the assumption that the brain finds the option that has the maximum probability of being correct, given incomplete and imperfect sensory information.

Earlier studies provide concrete evidence that humans and other animals perform probabilistic inference in a number of idealized situations. However, it is not always clear how these results translate to more realistic situations. In general, we can expect that Bayesian models describing human behavior can be fairly complex, and may require extensive and elaborate analytic computations. Often our observation time is short, and we need to make decisions in a short time. Further, most of our decisions involve processing of information from multiple sources. Given the constraints of biologically feasible computations, does our brain really make the best possible use of the information? If not, does it use some approximate strategies - and if so, when do such strategies fail? What are the possible models that best describe the computations performed by the brain?

To examine these questions, we study human behavior in an experimental setting (Chapter 2). The purpose of this study is to determine whether humans behave optimally in a fairly difficult perceptual task. We develop the theory for an

optimal Bayesian observer (Section 2.2), and provide several alternative suboptimal models (Chapter 4) that could possibly explain the observer's responses. These models encompass a range of assumptions about the observer's behavior. We compare these models (Chapter 5) using Bayesian and Akaike model comparisons (described in Chapter 3), to find the model that is most consistent with the experimental data.

Further, we theoretically analyze the performance of an ideal observer on a family of target detection tasks (Chapters 6 and 7). We present our analysis under certain assumptions about the parameters that determine the external structure of the task and those that govern the structure of observer's measurements.

In this Chapter, we give an overview of the key concepts of Bayesian inference. We begin by describing the different components of Bayes' rule, and how they can be interpreted in a perceptual task. We then describe the fundamentals of Bayesian modeling and their applications to psychophysical studies. Further, we provide the detailed explanation of Bayesian modeling using a simple target detection example. We conclude with a summary of the work presented in subsequent chapters of the dissertation.

1.1 Perception as Bayesian inference

Several experimental studies have provided evidence for Bayesian inference in perception [53, 43, 68, 84, 10, 83, 54, 94, 151, 99, 75, 150]. These studies also suggest that humans are capable of optimally integrating the available information. For instance, the experiments performed by Ernst and Banks [43] have provided evidence that human observers can optimally combine visual and haptic information to make a decision about the height of a raised object. More recently, several studies have established that humans are near optimal in finding a target among distractors [94, 98, 99]. The Bayesian framework presented in [78] indicates that humans not only integrate information based on the content present in the stimuli, but also based on the relevance of the task. These and many other similar studies are designed to investigate whether human behavior can be described in probabilistic terms.

Probabilistic computations can be helpful in routine life activities. For example, we try to predict the possibility of rain on a cloudy day based on weather conditions. We use information of ongoing vehicle and pedestrian traffic to make a decision about crossing a busy road. These activities require us to integrate different sensory information and make an informed decision. Generally, there is more than one possible choice or decision. Evidence suggests that we assign probability to the different options and base our decisions on them [83, 43, 93, 53]. According to Bayesian theory, an optimal observer computes probabilities for each possible event given the available information and makes a decision based on the most

probable state of the world. The theory relies on the assumption that the observer achieves this using Bayes' computation [12, 88].

1.1.1 Bayes' theorem

In the Bayesian framework, Bayes' theorem is used to compute the subjective belief about the state of the world based on accumulated evidence. The theorem depends on the computation of *conditional probabilities*. Conditional probabilities reflect the directional dependence between two events. For example, if A corresponds to the event of rain on a particular day, and B to the presence of clouds in the sky, then $P(A|B)$ indicates the probability that it will rain *given* a cloudy sky. We note that in general, $P(A|B) \neq P(B|A)$, and the two probabilities have different interpretations.

The Bayesian theory relies on the assumption that the brain combines the sensory measurements with our prior belief of the world via Bayes' formula,

$$\begin{aligned} P(\text{world state}|\text{data}) &= \frac{P(\text{data}|\text{world state})P(\text{world state})}{P(\text{data})} \\ &= \frac{P(\text{data}|\text{world state})P(\text{world state})}{\sum_k^{\text{Total states}} P(\text{data}|k^{\text{th}}\text{world state})P(k^{\text{th}}\text{world state})}. \end{aligned} \quad (1.1)$$

According to the Bayesian model of perception, we infer the probability of the world state given our sensory information, $P(\text{world state}|\text{data})$. We do so by using the probability of making a sensory measurement given a particular state of the world, $P(\text{data}|\text{world state})$. Importantly, we assume that this second probability is known to the observer and is part of their of perceptions model thereof.

The denominator in Eq. (1.1) is a normalization constant and ensures that the sum of the posterior probabilities of different world states is one. Frequently, we are interested in the expression in the numerator which establishes a direct proportional relation of the inferred probability to the sensory measurements and the associated prior belief. Thus, we often consider the following unnormalized version of Bayes' equation,

$$P(\text{world state}|\text{data}) \propto P(\text{data}|\text{world state})P(\text{world state}). \quad (1.2)$$

We note that given a particular observation, here "world state" is the variable and "data" is a constant. This formula combines our prior belief with available evidence to infer the state of the world. In the following section, we provide an intuitive interpretation of each term in the Bayes' formula. We follow the ideas and examples presented in [93].

1.2 Elements of Bayes' theorem

The Bayes' theorem consists of the following components: the likelihood function, the prior probability distribution, and the posterior probability distribution function.

The likelihood Function

Formally, $P(\text{data}|\text{world state})$ is written as $L(\text{world state}|\text{data})$ and is known as the *likelihood function* over possible world states given the sensory data. We note

1.2. ELEMENTS OF BAYES' THEOREM

that this is a function of the world state. It is not a probability distribution over the world states. Instead, it represents the likelihood of each state of the world given available data and summarizes the degree to which the sensory data favor one world state interpretation over the other. The accuracy of this selection is dependent on the quality of the observation.

The shape and nature of the likelihood function depend on the quality of the sensory input. The function has a clearly defined peak in case of high quality sensory input. It is usually flat or has multiple peaks when the information is inadequate or ambiguous. As an example, consider that we are walking on a street on a foggy day. Our visual information about incoming vehicle will be less accurate, and informative given unclear visible conditions. However, the sound and moving vehicle noise can be informative in such a case. Therefore, our likelihood function here will be concentrated around the world state favored by our auditory sensory information rather than visual one.

In principle, there are many factors that could affect the likelihood function. Weather, distance, and other physical conditions can influence the sensory measurements, hence the shape of the likelihood function. Moreover, the quality of sensory measurement varies across observers. Visual, auditory, and other sensory capabilities can impact the quality of an observer's sensory measurement, and hence the likelihood function.

The prior probability distribution

Prior knowledge about the world, $P(\text{world state})$, has significant impact on our inference process. It represents our beliefs or expectations about the world and determines the probability of each hypothesized state. Similar to the likelihood function, prior probability varies over hypothesized world states. Prior knowledge can evolve over time as we gather new information. Prior belief is subjective, as it depends on the observer's experiences. Hence, each observer may have a different prior distribution.

In our previous example, we can use our prior knowledge about the street and traffic conditions there while walking under unclear visible conditions. If the street has a curve or possible diversions, we can combine the prior information about such conditions with our current sensory observation while walking.

The posterior probability distribution

Decisions and inferences are based on the *posterior probability* of the world state, denoted by $P(\text{world state}|\text{data})$. The *posterior probability distribution function* represents the probability of each possible world state given our observational data and prior beliefs. To compute this probability, we use the likelihood function described above. Since the posterior probability is a combination of the likelihood function and prior information, the nature of the posterior distribution depends on these two factors. A sharp and peaked likelihood function results in a peaked and informative posterior probability distribution. On the other hand, in case of

flat likelihoods, posterior distribution resembles the prior distribution function and the observer does not gain any new information from sensory data.

Continuing with our example of walking on a street under foggy conditions, the likelihood function of a car approaching us could have a broader shape since our visual information is less informative. In such a case, a prior information about the traffic conditions on the street could be helpful in making a decision.

1.3 Bayesian modeling of perception

In this section, we discuss in detail the mathematical modeling of perception using Bayesian framework. Bayesian methods have been used to describe the process of perceptual inference and explain decisions of humans and animals on simple tasks [43, 94, 151, 68, 10, 83, 84]. Bayesian theory assumes that humans use Bayesian inference to update their belief about the state of the world. They do so by updating the posterior probability based on new sensory information. Also, this inference process is continuous and iterative since we incorporate our current state belief as prior information in making a new decision. There are numerous examples of routine life activities that can be explained using Bayesian inference. We discuss some of them below.

1.3.1 Visual and auditory perception

Visual and auditory perception are commonly studied examples of Bayesian inference. Recent experimental studies have focused on understanding visual perception through Bayesian modeling [43, 10, 125, 78, 144, 94, 151, 99, 37, 41, 42, 111, 7, 152]. These studies examine human behavior on simple visual tasks. We discuss some examples described further in [93].

Visual perception is of utmost importance to humans. We process multiple pieces of information contained in visual scenes to make decisions about the state of the world. One example is recognizing a friend in a crowd. If an observer is trying to find a friend in a large crowd from a distance, the visual information will have some degree of uncertainty. An ideal observer would compute a likelihood that each person in the crowd is the friend. As the observer gets closer to the friend, the quality of the sensory data improves and the likelihood function gets more peaked around the friend. Further, the observer uses some prior knowledge about the friend - for example, if the friend likes to wear black, then people wearing black will be assigned larger prior probability.

Our brain also handles tremendous amount of auditory information every day. We are exposed to numerous types of sounds, noise, and music in our environment, such as, music play, phone ring, alarm sound, vehicle horn, human speech, etc. Similar to visual perception, our auditory perception can be described as a process of Bayesian inference. For instance, when a song is played, we may try to guess its name. An ideal observer would compute a likelihood function over all

known songs. The observer would also use prior information about the songs or the type of music played by a radio station. In that case, the posterior distribution will be concentrated on frequently played songs.

These examples show how Bayesian computations could explain decisions in our daily routines. We now discuss how to design experiments to investigate whether human observers do employ Bayesian computations and how we model the collected data on the experiment. Here we give a brief description of *psychophysical tasks* that are extensively used to study perceptual behavior. We also illustrate the mathematical steps involved in the Bayesian modeling of such perceptual tasks.

1.3.2 Psychophysical studies

Psychophysical studies are used to analyze how animals process information from the physical world. These tasks are frequently designed to understand how animals integrate information contained in the stimuli to make a decision. The difficulty of these tasks usually depends on the characteristics of the presented stimuli. For example, an observer may be asked to discriminate whether a line is to the left or right of vertical. If the line is really close to vertical, the task can be difficult.

Observers are usually required to perform many iterations (*trials*) of the task. The trials can differ from each other if the characteristics of the stimuli are randomly chosen on each trial. The recorded responses of the observer can then be analyzed using Bayesian models. A wide range of psychophysical experiments

1.3. BAYESIAN MODELING OF PERCEPTION

have been analyzed using Bayesian inference approach [83, 84, 10, 43, 150, 53, 54, 98, 78, 144, 94, 68, 151, 75, 99] .

Bayesian modeling of perception consists of the following three steps:

1. describing the generative model,
2. specifying the inference process, and
3. computing the observer's estimate distributions.

These steps are the structural components of any perceptual Bayesian model and characterize the behavior of a model observer or a *subject* on the task. Below, we discuss each step in detail.

1.3.3 Step 1: the generative model

The generative model is a probabilistic model that describes the generation of the observer's sensory data. It mathematically describes the complete structure of the task. It is a forward directed graph with nodes representing the random variables characterized in the design of the task and the directed edges indicate the dependencies between variables. Each node has an associated probability distribution and the directed edge determines the influence of one variable on another, which is expressed in terms of *conditional probabilities*. At least one node in the model corresponds to the variable of interest describing the state of the world and another variable is the observer's sensory data. The observer infers the latent variable of the world state from the stimulus shown in the task.

We denote the feature or characteristic of a stimulus by s . The feature is sometimes itself referred to as the stimulus. Depending on the problem, the variable of interest, which we denote by W , could be different from the stimulus, s itself. It is frequently assumed that the observer makes a noisy measurement, x , of the presented stimulus. The measurement is also sometimes referred as the *internal representation* or the observation of the stimulus. Thus, the generative model contains the variable of interest (if different from the stimulus), the stimulus, and its measurement. Figure 1.1(A) shows the graphical representation of the generative model and the probability distributions associated with each node.

1.3.3.1 Distributions in the generative model

The probability distributions in the generative model can be determined from the experimental design of the task. The *world state distribution* or prior distribution, denoted by $p(W)$, represents the distribution of probabilities over the world states in absence of any sensory information. This distribution could either be discrete or continuous depending on the associated random variable. The *stimulus distribution* is a function of the world state variable, W , and is denoted by $p(s|W)$. This distribution is completely specified by the design of the experiment. When a stimulus itself represents the state of the world, we have $p(s|W) = p(s)$.

Measurements or sensory data are usually noisy. This noise could come from many sources: the random variability due to intrinsic stochastic processes, limitations of our sensing capability, and other unknown sources. We need to make an assumption about the noise in the generative model. Even in response to the




Figure 1.1: **Bayesian modeling of a perceptual task - defining the generative model and deriving the inference process.** **(A) The generative model.** The first step in Bayesian modeling is to define the generative model. This figure outlines the graphical representation of the generative model we will be using throughout the dissertation. The nodes represent the variables involved in the task and arrows determine the influence of one node on another. This influence is mathematically described in terms of conditional probabilities. The observer infers the (hidden) state of the world, W from the stimulus, s presented in the task by making a measurement, x of the stimulus. **(B) Inference process.** The second step in Bayesian modeling is to derive the inference process of an observer. That is, to understand the mathematical process by which the observer infers W based on the measurement, x . This step involves inverting the generative model and marginalizing over intermediate variable, s to compute a decision criterion and making an estimate, \hat{W} about the state of the world.

1.3. BAYESIAN MODELING OF PERCEPTION

same stimulus, the observer's measurements vary randomly over the course of experimental trials. The distribution of the measurement, x , given the stimulus, s , is a conditional distribution which is denoted by $p(x|s)$. It represents the probability with which a stimulus results in a measurement, x . Frequently, we use the Gaussian distribution to model measurement noise,

$$p(x|s) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-s)^2}{2\sigma^2}}. \quad (1.3)$$

The standard deviation, σ , of the Gaussian function reflects the uncertainty or noise in the measurement. A higher (lower) value of σ reflects the low (high) quality of the measurement and is associated with a wider (narrower) measurement distribution. The inverse of the variance, $\frac{1}{\sigma^2}$, is commonly known as the *precision* or *reliability* of the measurement.

1.3.3.2 Used prior distributions in the experiment

The world state distribution reflects an observer's prior belief about the state of the world. Subjects can make incorrect assumptions about the prior state of the world. Given the set up of a psychophysics experiment, it might be difficult for an observer to correctly determine the world state distribution. Subjects could use a prior based on the experiences in natural world, but this could potentially be very different from the experimental world state distribution. For example, subjects could have a higher prior probability for vertically and horizontally aligned objects based on their experiences whereas the objects in an experiment could possess any orientation. Also, a prior of light coming from above would be stronger

than any other direction for similar reasons. Priors based on common experiences are more likely to have higher probability. Therefore, sensible assumptions need to be made about a subject's prior in the experiment. In most practical cases, a flat or uniform prior is assumed that assigns equal probability to each outcome for the state of the world, that is, $p(W) = \text{constant}$. Such a prior simplifies the inference computations. Alternatively, we can also determine the subject's prior from experimental data. This practice is commonly used in the case of binary world state variables.

In summary, the distributions $p(W)$, $p(s|W)$, and $p(x|s)$ completely define the generative model of the task and constitute a major component of Bayesian modeling of perceptual inference.

1.3.4 Step 2: the inference or perception process

In the next step of Bayesian modeling, we specify a model to determine how an ideal observer makes decisions. In our computations, we use certain assumptions about the observer's measurements (specified in the generative model). As discussed earlier, the inference process involves computation of posterior probability distribution given the likelihood function and the prior distribution. It essentially involves the "inversion" of the generative model in order to perform computation about the world state, W given the sensory data, x . Given the posterior distribution, denoted by $p(W|x)$, the observer makes a single estimate of the world state.

Commonly, the observer follows the *maximum-a-posteriori estimation (MAP)* to obtain the point estimate that has the highest posterior probability. MAP estimation is one of the most common ways of reading out the posterior distribution as it finds the most probable option. The prior distribution $p(W)$, the likelihood function $L(W|x)$ or $p(x|W)$, and the posterior distribution $p(W|x)$ are key components in the inference process. Figure 1.1(B) describes a general scheme of the inference model in the Bayesian modeling.

We can also use the likelihood function to make the best guess of the stimulus by maximizing the function over hypothesized world states. This estimate is called as the *maximum-likelihood estimate (MLE)* of W and is denoted by \hat{W}_{ML} . Mathematically, we write

$$\hat{W}_{\text{ML}} = \arg \max_W L(W|x). \quad (1.4)$$

1.3.4.1 Marginalization

We note that the observer is interested in determining W , and not s . Hence, the generative model includes a variable that is not of our interest, but it provides important information for computing the posterior distribution function. Such a situation is dealt with the *marginalization* process, where the information about the intermediate variable is averaged out. It is a commonly used technique in a Bayesian model, that integrates or sums the values of all such auxiliary variables to obtain the desired probability distribution over the parameter of interest. Although auxiliary variables, such as the stimulus, are not of primary interest,

they play a critical role in the generative model and must be accounted for via marginalization to obtain an accurate perception.

In the present case, we need to compute the posterior distribution $p(W|x)$, and not $p(s|x)$. We can obtain it using the Bayes' formula described in Eq. (1.2)

$$p(W|x) \propto p(x|W)p(W), \quad (1.5)$$

The distribution $p(x|W)$ is not specified in our Bayesian model, instead we have information about the noise distribution, $p(x|s)$. Thus, we marginalize the information over the intermediate variable, s to obtain the required distribution $p(x|W)$ as follows

$$\begin{aligned} p(x|W) &= \int p(x|s, W)p(s|W)ds, & \text{if } s \text{ is continuous,} \\ &= \sum_i p(x|s = s_i)p(s = s_i|W), & \text{if } s \text{ is discrete.} \end{aligned}$$

The above marginalization step, also shown in Figure 1.2 links the world state variable, W to the measurement, x via the intermediate variable, the stimulus, s .

We further note that the shape of the posterior distribution is preserved under the normalization constant in Eq. (1.5). After computing the posterior distribution, the observer then reads the maximum-a-posteriori estimate, denoted by \hat{W}_{MAP} , by maximizing the posterior distribution

$$\hat{W}_{\text{MAP}} = \arg \max_W p(W|x). \quad (1.6)$$

The MAP estimate, \hat{W}_{MAP} is also the mode of the posterior distribution and reflects the observer's estimate of the world state of interest.




Figure 1.2: **The process of marginalization in Bayesian modeling.** The generative model usually contains auxiliary variables that are not of primary interest, but they may have necessary information about the state of the world, W . Here the stimulus, s is an intermediate variable, but it links the world state variable of interest with the measurement, x . Marginalization is a process to deal with such ancillary variables to obtain the desired expression for the likelihood or posterior probability of the world state variable of interest. It involves averaging or integrating over the possible values of the ancillary variable and is very common in Bayesian modeling.

1.3.5 Step 3: the observer's estimate distribution

The observer's measurement, x heavily depends on the sensory noise. Even under same experimental conditions, the measurements vary across trials due to different sensory noise. Therefore, x is a random variable across experimental trials. As a result, the MAP estimate, \hat{W}_{MAP} is also a random variable in response to a fixed stimulus and has a probability distribution. In the final step of Bayesian modeling, we thus determine the distribution of the observer's estimates. We note that the mapping from measurement to MAP estimate is completely deterministic and the randomness in MAP estimate is only because of variability in the measurements from trial to trial.

Since we only have access to the MAP estimates of the observer, we need to compute the probability of each possible estimate in a particular experimental condition. We can then compare the predictions of the model with the observer’s behavior. Thus, we compute the distribution of MAP estimates given a fixed stimulus, say $s = s_{\text{true}}$. This is usually denoted by $p(\hat{W}_{\text{MAP}}|s_{\text{true}})$ and predicts how likely the estimate is given the fixed stimulus, s_{true} .

To summarize, Bayesian modeling of a perceptual task consists of specifying the generative model, deriving the inference process for the observer, and evaluating the distribution of MAP estimates over many trials. This mathematical tool of quantifying the perceptual behavior is schematically shown in Figure 1.3.



Figure 1.3: **Steps involved in Bayesian modeling of a perceptual task.** The figure presents the schematic of a Bayesian inference process to model a perceptual task. We will follow this plan for all the tasks discussed in the dissertation. The first step of specifying the generative model involves describing the probability distributions to understand how sensory data are generated from the state of the world. The observer makes an estimate of the world state based on the sensory measurement on each trial of the task. This constitutes the second step of deriving the inference process in a Bayesian model. The estimate of the observer varies across trials in response to a fixed stimulus and follows a distribution. In the final step of Bayesian modeling of the task, this estimate distribution is computed.

Thus far, we have described the general process of modeling a perceptual task using a Bayesian approach. In the dissertation, we use Bayesian models to understand visual perceptual inference. In the following section, we discuss a particular

family of vision based perceptual tasks, known as *visual search tasks*. Specifically, we elaborate the process of Bayesian modeling in a simple example of such a task and discuss the possible questions which need to be explored further. We will examine those questions in extensive details in subsequent chapters.

1.4 Visual search

Visual search is a common example of vision based perceptual task. This involves an active scan of multiple objects for a particular object or feature of interest, referred as *the target* among other objects or features, *the distractors*. Finding a friend in a large crowd or finding a particular set of keys among other similar items, or locating an insect hidden in the corner are some examples of visual search from our everyday life. These examples also highlight the importance of performing visual search in our normal life. But, doing psychophysics with natural scenes is challenging. The natural scenes are high-dimensional and highly structured. They are so rich in content that a single mathematical model may not be plausible to capture all the characteristics of a scene. Moreover, noise in natural scenes is largely unknown and perhaps has a complex correlation structure that is harder to capture with simple probability distributions. Further, a distinct object classification may be unavailable in case of natural scenes. Therefore, for the purposes of psychophysical studies conducted in the laboratory, quite simplified visual search tasks are considered. These psychophysical tasks usually contain some highly distinct objects that only differ along a small number of stimulus dimensions. Clearly,

these simple and fabricated tasks do not replicate natural scenes; however, they serve as a practical tool for understanding the perceptual inference computation performed by the brain.

The ability to consciously locate an object (target) among a complex array of stimuli (distractors) has been extensively studied in psychophysics over many years [108, 37, 148, 43, 41, 8, 42, 111, 112, 7, 152, 94, 98, 99]. These studies also validate the modeling of our perceptual behavior using Bayesian approach.

1.4.1 An example of a target detection task

We now discuss a specific example of visual search task, namely, a *target detection task*. Our example is similar to the one discussed by Ma et al. in [93]. We consider a simple task with only two stimuli. The observer needs to report whether a *target* stimulus is present in the scene. Stimulus orientation is the task-relevant feature. We elaborate the mathematical steps involved in the Bayesian modeling of the task.

1.4.1.1 Step 1: Generative model

The observer is presented with two stimuli on a visual display. These stimuli could either be bars, gratings, or ellipses characterized by their orientation or eccentricity. A target is a stimulus with a particular characteristic. We assume that a target is a stimulus with vertical orientation, denote its orientation by $s_T = 0$, and measure stimulus orientation relative that of a target. A *distractor* is defined as the

stimulus having a non-target orientation. The observer reports whether a target is present in the visual display on each trial. We denote the target presence by a binary variable, T , so that $T = 1$, if a target is present and $T = 0$ if absent. The target is present with a probability of 0.5 in each trial and it can be present at either of the two locations. Thus, the location of the target is unknown to the observer and hence, the observer needs to take this into account by marginalizing over both possible target locations. Figure 1.4(C) shows the three possible displays in the task. We also represent the target presence at location, i , by a binary variable, T_i and further define the spatial location vector as $\mathbf{T} = (T_1, T_2)$. We denote the stimulus orientation of the i^{th} stimulus by s_i .

When the target is absent, *i.e.*, $T = 0$, each stimulus orientation is drawn from a normal distribution with mean 0 and standard deviation, σ_s . We write (see notation in Appendix A)

$$s_i | T = 0 \sim \mathcal{N}(0, \sigma_s^2). \quad (1.7)$$

We denote the probability density function of the normal distribution, $\mathcal{N}(0, \sigma_s^2)$ by $f(s_i; 0, \sigma_s^2)$, where

$$f(s_i; 0, \sigma_s^2) = \frac{1}{\sqrt{2\pi\sigma_s^2}} \exp\left(-\frac{s_i^2}{2\sigma_s^2}\right).$$

The assumption of the Gaussian noise is reasonable since the stimuli are not placed too close to each other on the screen. When $T = 1$, one of the stimuli is chosen as the target with uniform probability. If the target is present at location, j , for $j \in \{1, 2\}$, then $s_j = s_T = 0$ and we choose the distractor orientation according to Eq. (1.7). We assume that the observer makes independent (between locations)

and noisy measurement of the stimulus, s_i , denoted by x_i , which is drawn from the following normal distribution

$$x_i | s_i \sim \mathcal{N}(s_i, \sigma_i^2). \quad (1.8)$$

The noise, σ_i determines the uncertainty in the i^{th} measurement of the stimulus, s_i and is known to the observer. The generative model of the task is illustrated in Figure 1.4(A).

1.4.1.2 Step 2: Inference

The optimal Bayesian observer infers target presence based on the stimulus measurements and using the information about the generative model. The observer computes the *log posterior ratio (LPR)* of target presence from the measurements, x_1 and x_2 as

$$d(x_1, x_2) = \log \frac{p(T = 1 | x_1, x_2)}{p(T = 0 | x_1, x_2)}, \quad (1.9)$$

and reports "target present" if $d > 0$ and "target absent" otherwise. The variable, d is known as the Bayesian decision variable. Using the Bayes' formula, we rewrite the above equation as the sum of *log likelihood ratio (LLR)* and *log prior ratio*,

$$d(x_1, x_2) = \log \frac{p(x_1, x_2 | T = 1)}{p(x_1, x_2 | T = 0)} + \log \frac{p(T = 1)}{p(T = 0)}.$$

Since the target is present or absent with equal probability, we have a uniform prior over T , and therefore,

$$d(x_1, x_2) = \log \frac{p(x_1, x_2 | T = 1)}{p(x_1, x_2 | T = 0)}. \quad (1.10)$$




Figure 1.4: **Bayesian modeling of a simple target detection task with two stimuli.** (A) **The generative model.** The binary variable, T describes the target presence in a trial. The two stimuli, s_1 and s_2 are chosen conditioned on T . When $T = 1$, one of the stimuli is a target with a vertical orientation, while the orientation of the other stimulus is chosen randomly from a normal distribution. The observer makes noisy and independent measurements, x_1 and x_2 of the two stimulus. (B) **The inference process.** The observer combines the two measurements to compute a decision variable, $d(x_1, x_2)$ and infers an estimate, \hat{T} of the world state variable, T . The decision variable, $d(x_1, x_2)$ is a log posterior ratio of the probability of reporting "target present" and "target absent", given the observer's measurements. If $d > 0$, the observer reports target is present and absent otherwise. (C) **Example displays in the task.** Since there are only two stimuli and one target, three types of visual displays can be presented to the observer. In the first two displays, the target is present to the left and right of the cross in the center. When there is no target, both stimuli are distractors and have randomly chosen orientations. The bottom display illustrates such an example.

We compute the numerator in Eq. (1.10) by marginalizing over the two possible target locations and stimuli, s_i in the following equations

$$\begin{aligned}
 p(x_1, x_2 | T = 1) &= \sum_{j=1}^2 p(x_1, x_2 | T_j = 1, T = 1) p(T_j = 1 | T = 1) \\
 &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^2 p(x_j | T_j = 1) p(x_i | T_i = 0) \\
 &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^2 \int p(x_j | s_j) p(s_j | T_j = 1) p(x_i | s_i) p(s_i | T_i = 0) ds_j ds_i \\
 &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^2 \int f(x_j; s_j, \sigma_j^2) \delta(s_j - 0) f(x_i; s_i, \sigma_i^2) f(s_i; 0, \sigma_s^2) ds_j ds_i \\
 &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^2 f(x_j; 0, \sigma_j^2) f(x_i; 0, \sigma_i^2 + \sigma_s^2) \quad (\text{using Eqs. (B.1) and (B.4)}) \\
 &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^2 \frac{1}{2\pi\sqrt{\sigma_j^2(\sigma_i^2 + \sigma_s^2)}} \exp\left(-\frac{x_j^2}{2\sigma_j^2} - \frac{x_i^2}{2(\sigma_i^2 + \sigma_s^2)}\right).
 \end{aligned}$$

When the target is absent, both stimuli are distractors. Thus, we compute the denominator in Eq. (1.10) as

$$\begin{aligned}
 p(x_1, x_2 | T = 0) &= p(x_1 | T_1 = 0) p(T_1 = 0 | T = 0) p(x_2 | T_2 = 0) p(T_2 = 0 | T = 0) \\
 &= \int p(x_1 | s_1) p(s_1 | T_1 = 0) p(x_2 | s_2) p(s_2 | T_2 = 0) ds_1 ds_2 \\
 &= \int \prod_{i=1}^2 \left(f(x_i; s_i, \sigma_i^2) f(s_i; 0, \sigma_s^2) ds_i \right) \\
 &= \prod_{i=1}^2 f(x_i; 0, \sigma_i^2 + \sigma_s^2).
 \end{aligned}$$

We now substitute the above computed individual likelihoods for target present, and absent cases to obtain an expression for Bayesian decision variable, d :

$$\begin{aligned}
 d(x_1, x_2) &= \log \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^2 \frac{f(x_j; 0, \sigma_j^2) f(x_i; 0, \sigma_i^2 + \sigma_s^2)}{f(x_j; 0, \sigma_j^2 + \sigma_s^2) f(x_i; 0, \sigma_i^2 + \sigma_s^2)} \\
 &= \log \frac{1}{2} \sum_{j=1}^2 \frac{f(x_j; 0, \sigma_j^2)}{f(x_j; 0, \sigma_j^2 + \sigma_s^2)} \\
 &= \log \left[\frac{1}{2} \sum_{j=1}^2 \left(\frac{\sigma_j^2 + \sigma_s^2}{\sigma_j^2} \right)^{1/2} \exp \left(\underbrace{-\frac{x_j^2}{2\sigma_j^2}}_I - \underbrace{\frac{x_j^2}{2(\sigma_j^2 + \sigma_s^2)}}_{II} \right) \right]. \quad (1.11)
 \end{aligned}$$

The above expression gives us the decision variable on the task. An ideal Bayesian observer performing the above described target detection task makes a decision based on this decision variable. The decision variable, $d(x_1, x_2)$ depends on the precision of the measurement and also on the external variability of the distractor, σ_s^2 . Each exponent term in the above expression provides an evidence towards j^{th} stimulus being a target: (I) if the j^{th} measurement is close to the vertical orientation, this term corresponds to an increased likelihood for the j^{th} stimulus being the target, while (II) the second term decreases such a likelihood. The appropriate scaling of the two measurements by inverse of the external and internal noise variances determine the correct likelihood for the j^{th} stimulus being the target or a distractor.

1.4.1.3 Step 3: MAP estimate distribution

We denote the observer's MAP estimate of T by \hat{T} . When $d > 0$, the MAP estimation is to report "target present" and "target absent" otherwise. The distribution of the MAP estimate describes the behavior of the Bayesian observer across many trials and involves computing the probabilities that the observer will report "target present" when the target is actually present and when it is absent. That is, we need to compute $p(\hat{T} = 1|T = 1)$ and $p(\hat{T} = 1|T = 0)$. These probabilities are also known as the *hit* and *false-alarm rates*, and are computed given a fixed set of stimulus, s_1 and s_2 :

$$\begin{aligned} p(\hat{T}|s_1, s_2) &= \int p(\hat{T}|x_1, x_2)p(x_1|s_1)p(x_2|s_2)dx_1dx_2 \\ &= \int \delta_{\hat{T}, \text{sgn}(d(x_1, x_2))}p(x_1|s_1)p(x_2|s_2)dx_1dx_2. \end{aligned}$$

Here δ represents the Kronecker delta function. As the decision variable, d computed in Eq. (1.11) is a non-linear function of x_1 and x_2 , the above integral is analytically intractable, and hence needs to be approximated using Monte Carlo simulations in practice. This completes the final step in the Bayesian modeling of the above described target detection task.

1.4.1.4 A suboptimal model

Eq. (1.11) describes the decision variable for an optimal Bayesian observer to perform the task. However, it is not necessary that the observer will follow this rule. Specifically, it is possible that an observer may not be optimal in making a decision on the task and may use some other decision strategy to make a decision.

The decision variable, $d(x_1, x_2)$ computed in Eq. (1.11) is clearly non-linear and implicitly depends on other parameters that characterize the structure of the task. For instance, the external variance of the distractor stimulus, σ_s^2 and the noise in making a measurement, σ_i^2 clearly affect this variable. We generally assume that the observer is aware of the internal noise with which the measurement is made, but he might not know the external variance that determines the structure of the task. In such a case, the observer would use an incorrect assumption about the generative model in making a decision. Further, the observer could also make a guess on each trial without using any information about the task.

Therefore, the following questions frequently emerge in analyzing the responses of the observer on such tasks: What model does the observer follow in making a decision on the task? How do we infer the parameters of a model that fit the subject's data? And in the case of multiple models, how can we compare models to select the one that best describes the experimental data? We will examine these questions in great length in Chapters 3 to 5 and present detailed analysis for a particular target detection task described in Chapter 2.

We now consider a suboptimal model to understand the observer's behavior on the target detection task. Let us assume that the observer is not optimal and instead makes a decision based on the minimum of the two measurements of the stimuli. If the minimum of the measurements is below some threshold, θ , the observer reports "target present", and absent otherwise. We denote the decision variable for such a model by $d_{\min}(x_1, x_2)$ and define

$$d_{\min}(x_1, x_2) = \min(|x_1|, |x_2|). \quad (1.12)$$

The observer compares whether $d_{\min}(x_1, x_2) < \theta$ to make an estimate $\hat{T} = 1$. We note that the observer's decisions based on the decision variable, $d_{\min}(x_1, x_2)$ are suboptimal. The threshold models have been used in many earlier studies to model subjects' responses on psychophysical tasks [112, 152, 8, 42, 55, 94].

Figure 1.5 compares the performance of an observer as a function of the standard deviation of the stimulus distribution, σ_s , for the optimal and threshold models. Here we assume that the observer uses equal precision for both measurements, that is, $\sigma_1^2 = \sigma_2^2 = \sigma^2$. We consider the performance of the observer based on the following models: (i) optimal model, (ii) threshold model with $\theta = \frac{1}{2}\sigma$, (iii) threshold model with $\theta = \sigma$, and (iv) threshold model with $\theta = \frac{3}{2}\sigma$. Thus, we consider threshold models where the parameter θ depends on the precision of the observers' measurements.

We note that the performance increases for all models as the standard deviation of distractor orientations, σ_s increases. This is expected since the task becomes easier as the orientations of the distractors deviates away from that of the target. Also, we observe that the performance predicted by threshold models is lower as compared to the optimal model when the threshold parameter, θ is small. However, for a sufficiently large threshold value (here $\theta = \frac{3}{2}\sigma$), the observer behaves similarly according to both optimal and threshold models at large values of σ_s . We find that all threshold models closely predict the performance as that of the optimal model at low values of standard deviation between distractors. At low values of σ_s , the distractors are more likely to have orientations close to that of the target




Figure 1.5: **Performance of an observer in a simple target detection example based on different models.** Proportion of correct responses as a function of the standard deviation of distractor orientations, σ_s for an observer based on different models. The optimal model has the maximum performance at all values of σ_s than other threshold models. A lower performance is observed at low standard deviations for all models since it becomes difficult to detect a target on the task among distractors that have relatively similar orientations to that of the target. As the standard deviation of the distractor orientations increases, the task becomes relatively easier and the performance increases for all models. The model with high threshold parameter, θ predicts a similar performance of the observer as that of the optimal model. This indicates that it is difficult to choose a model that is most consistent to describe the observer's behavior.

(see Eq. (1.7)). In such a case, it will be difficult for the observer to make a decision and will have lower performance as predicted by both optimal and threshold models.

Therefore, the decisions of the observer depend on the model parameters and the precision of the measurements. In order to understand the responses of the observer, we need to estimate these model parameters and also determine the precision level of observer's measurements. Using these parameter estimates, we make predictions for the observer's responses based on different plausible models. These models can be very close in their predictions (for example, optimal and threshold model with $\theta = \frac{3}{2}\sigma$ in Figure 1.5) and it may be difficult to find the model that is most consistent with the observer's behavior.

1.4.2 Generalizations

In general, visual search task comprises of group of tasks: *target detection* - determining whether a target is present or not in a scene; *target localization* - finding the location of the target when the target is always present and *target discrimination or classification* - classifying the target to one of the pre-defined categories. We only focus on the target detection tasks in the dissertation and explore different parameter relations in these tasks.

The target detection task described in Section 1.4.1 is an extremely simple example of a target (or visual) search task with only two stimuli. However, in general, we make decisions in presence of a large number of distractors. For instance,

1.4. VISUAL SEARCH

in order to find a friend in a large crowd, we need to carefully scan each individual and undergo an identification process. Also, the distractor orientations may not always be independent. They could possess an unknown complex structure and we must account for such structures to make correct decisions.

Several research studies have considered the visual search tasks with a reasonably large number of stimuli [8, 75, 44, 111, 98, 135, 42, 99, 151, 94, 150, 137]. Some of them have focused on analyzing performance as a function of set-size [111, 8, 42, 135, 44, 136, 98, 99]. Furthermore, the experimental studies done by Ma et al. [94, 151, 98, 99] have explored the behavior of the subjects on a search task with two types of distractors: when all distractors have identical orientation and when the distractors possess different independent orientations.

In addition to the possibility of varying set size and structural orientations of the distractors, the task could also have multiple targets in the visual display. Such a possibility has not been explored in scientific studies. We thus examine this possibility with complete mathematical details in Chapter 7 and further analyze the impact of different parameter correlations on the performance of an optimal observer on the task. In real life, there are several examples where multiple targets are present and our brain needs to process the information to find at least one of them. For instance, we may need to find a blue marker in a box of black and blue markers. There could be more than one blue marker present in the box among black markers that serve as distractors.

Thus far, we have described the principles of Bayesian approach used in building perceptual models. These concepts are fundamentals for the work presented

in the dissertation. We now present the outline of the dissertation and the main questions examined in the following chapters.

1.5 Outline of the dissertation

The dissertation is split into two parts. In the first part, we study how humans make decisions in response to structured input. In the second part, we examine how structured measurements affect the decisions along with structured stimulus input.

Visual search for a single target among distractors, with a single relevant feature has been studied extensively. These studies have largely focused on two types of distractors: distractors with identical orientations and with independent random orientations. Therefore, either distractors are exactly alike or they differ from each other across trials. In a target-present trial with identical distractors, the target would be an odd-ball stimulus, and hence can be detected easily. However, if the distractors have independent orientations, there would be hardly any structure in the scene that could possibly help the observer in finding a target. The experimental studies [94, 99] showed that humans decisions are consistent with optimal Bayesian models in case of both types of distractors. But, these two distractor conditions represent the extreme structural possibilities: from high structure to none. However, natural scenes possess more complex structure; and the objects in the scenes can be correlated with each other in many possible ways. Therefore, this raises the question how humans make decisions in response to

1.5. OUTLINE OF THE DISSERTATION

weakly structured input? In particular, do humans take into account the weak stimulus correlations and do they make near-optimal decisions in such a case?

We examine these questions in Chapters 2 to 5. We consider a target detection task with N stimuli and one target. The distractors are assigned orientations that have different amount of correlations. The varying amount of correlations among distractor orientations allows us to introduce structure in the visual scene. The structure of these correlations must be taken into account to make optimal decisions about target presence. In Chapter 2, we derive the mathematical theory for a Bayesian optimal observer. We also performed a psychophysical experiment based on the design of the task and analyze the collected human subjects' data in Chapters 4 and 5.

We explore whether humans are optimal in inferring correct correlation strength among distractor orientations. We test several Bayesian models that could possibly explain subjects' behavior on the experiment and fit them to the data. The fitting of a model requires finding the model parameters that provides best fit to the subjects' responses. Further, we need model selection techniques to find the best one. We describe the maximum-likelihood parameter estimation and model comparison techniques in details in Chapter 3. We use these techniques to fit different model parameters in Chapter 4 and select the best model for the data in Chapter 5.

For the purposes of our analysis in Chapters 2 and 5, we have assumed the noise in the measurements to be independent and normally distributed. Both the

1.5. *OUTLINE OF THE DISSERTATION*

assumptions about independence and Gaussianity can be questioned. In particular, there is some evidence that neural correlations can be present at long distances in visual field suggesting that sensory measurements will be correlated [39, 31, 30, 124]. However, much of the visual search studies lack the assumption about measurement correlations in the sensory measurements. We thus focus on the effects of such an assumption on the performance in categorical and global perceptual judgements. To make a correct decision in this case, the observer needs to take into account not only the correlations between the measurements, but also the statistical structure of the stimuli.

In the second part of the dissertation (Chapters 6 and 7), we explore the joint effects of measurement and stimulus correlations in a family of visual search tasks. To investigate how the interaction between both correlations, stimulus and measurement, can affect the decisions of an ideal observer, we consider the assumption of correlated sensory measurements in the target detection task described in Chapter 2. Thus, in Chapter 6, we study a target detection task with a single target and assume that the sensory measurements are correlated following a multivariate normal distribution. We provide complete details of the analytical computations that an ideal observer follows to make correct decisions on the task. We analyze the impact of statistical structure of the scene along with measurement correlations on the performance of the ideal observer.

Further, in Chapter 7, we continue with our examination and analysis in the case of multiple targets in the detection task. Multiple targets having identical orientation would introduce more statistical structure in the scene as compared

1.5. OUTLINE OF THE DISSERTATION

to a single target. Even if the distractors possess independent orientations, having multiple targets increase the chance of detecting a target in a pool of stimuli. Therefore, we inspect how external statistical structure present interacts with the structure of the measurements to affect the decisions of the observer. In particular, we analytically and numerically analyzed how the performance of the ideal observer behaves as a function of different parameters that determine the structure in the external scene and in the sensory measurements. We find that the performance changes considerably in the case of multiple targets, while it remains unchanged when only a single target is present. Therefore, the accuracy of decisions on these visual search tasks are greatly influenced by the relationship between the stimulus and measurement correlations.

Stimulus correlations in a target detection task

Bayesian models have been successfully used to study perceptual behavior. Many studies have used these quantitative models to show that humans perform near-optimally on simple perceptual tasks. That is, optimal Bayesian models successfully explained the behavior of human observers on these tasks. These models assume that observers make best possible decisions given the uncertain and noisy sensory measurements.

A number of recent studies have analyzed human behavior in simplified scenes containing multiple objects. Human behavior was found to be close to Bayes-optimal on different psychophysical tasks: visual search [94, 98, 99], sameness judgement [151], and change detection [75]. However, many of the visual search

studies have only considered two types of distractors: *homogeneous* and *heterogeneous* distractors [147, 37, 162, 105, 152, 125, 94, 98, 106, 99]. Homogeneous conditions have identical distractors, while heterogeneous conditions have independent, randomly oriented distractors. These studies considered the orientations of homogeneous distractors to be same across all trials, while they varied the orientation of distractors across experimental trials in the heterogeneous condition. Thus, orientations of the distractors in the homogeneous condition were predictable from trial to trial.

Recent studies done by Mazyar et al. [99] studied the human behavior on a target detection task under the violation of trial-to-trial predictability in the case of homogeneous distractors. They used Gaussian distributions to randomly draw the orientation of distractors in the case of homogeneous condition. They also performed experiments to study human decisions in response to heterogeneous distractors and found that humans were near-optimal in detecting a target among both types of distractors.

However, their studies were also limited to homogeneous and heterogeneous distractors. By contrast, visual stimuli in natural scenes possess a complex and higher-order structure. The orientations of the objects in natural scenes are correlated to different extent with each other. It is therefore important to examine how visual perception is affected by structured input. Specifically, to understand how differently correlated input affect our decisions.

We examine these questions in a psychophysical task. We study the decisions of human observers on a target detection experiment under the effect of structured

input. We introduce structure in visual scenes by varying the amount of correlations between distractor orientations. We note that homogeneous condition refers to perfectly correlated (identical) distractors, while heterogeneous distractors are uncorrelated. In our study, we interpolated between the heterogeneous and the homogeneous conditions, and studied the intermediate regimes of partial correlations. We are interested in understanding how human observers make decisions in response to differently correlated (uncorrelated, partially, and perfectly correlated) stimuli. The intermediate regime of correlations can be challenging from an observer's point of view, since the stimuli are only partially correlated, and only introduce a weak structure in the scenes. The observer needs to take into account the strength and structure of the stimulus correlations to make an optimal decision.

Therefore, we examine the following questions in our study: Do humans take into account the strength and structure of the stimulus correlations? And if they do, can they make near-optimal decisions? We provide answers to these questions for the target detection task in Chapter 4. We note that visual search is one particular example of a task where these questions are relevant. They are more generally applied to a variety of perceptual tasks.

We begin this chapter with the description of the model for the target detection task. We continue with the derivation of the inference process for an optimal Bayesian observer. An ideal observer makes decisions according to the derived decision rule to infer target presence on the task. Later in the chapter, we present the details of the task based experiments we have conducted.

2.1 Generative model

We consider the following target detection task: An observer is presented with N stimuli. The observer reports whether a vertical target stimulus is present or absent among a group of distractors. Stimuli are characterized by their orientations. We denote the target orientation by $s_T = 0$. We represent target presence by the binary variable, T , so that $T = 1$ if the target is present, and $T = 0$ if absent. This notation is consistent with studies done by Ma et al. [94, 98, 99]. Target presence at location i is similarly represented by a binary variable, T_i . We also denote the spatial location vector by $\mathbf{T} = (T_1, T_2, \dots, T_N)$. In each trial, the target is present with a $1/2$ probability.

We denote the orientations of the stimuli by $\mathbf{s} = (s_1, s_2, \dots, s_N)$. When $T = 0$, the target is absent and all stimuli are distractors. Therefore, $\mathbf{T} = (0, 0, \dots, 0)$ and we write

$$p(\mathbf{T}|T = 0) = \delta(\mathbf{T} - \mathbf{0}_N).$$

Here a subscript denotes the length of a vector, so that the vector $\mathbf{0}_N$ has N components. In this case, the orientations of the stimuli are drawn from an N -dimensional multivariate normal distribution with mean vector, $\mathbf{s}_D = (s_D, s_D, \dots, s_D)$ and covariance, Σ_s , and we write

$$\mathbf{s}|T = 0 \sim \mathcal{N}(\mathbf{s}_D, \Sigma_s). \quad (2.1)$$

2.1. GENERATIVE MODEL

We denote the probability density function of the multivariate normal distribution, $\mathcal{N}(\mathbf{s}_D, \Sigma_s)$ by $f(\mathbf{s}; \mathbf{s}_D, \Sigma_s)$, where

$$f(\mathbf{s}; \mathbf{s}_D, \Sigma_s) = \frac{1}{\sqrt{(2\pi)^N |\Sigma_s|}} \exp\left(-\frac{1}{2}(\mathbf{s} - \mathbf{s}_D)^T \Sigma_s^{-1} (\mathbf{s} - \mathbf{s}_D)\right). \quad (2.2)$$

The $N \times N$ covariance matrix, Σ_s , contains identical diagonal entries, σ_s^2 (variances) and identical off-diagonal entries, $\rho_s \sigma_s^2$ (covariances):

$$\Sigma_s = \begin{bmatrix} \sigma_s^2 & \rho_s \sigma_s^2 & \cdots & \rho_s \sigma_s^2 \\ \rho_s \sigma_s^2 & \sigma_s^2 & \cdots & \rho_s \sigma_s^2 \\ \vdots & & \ddots & \vdots \\ \rho_s \sigma_s^2 & \rho_s \sigma_s^2 & \cdots & \sigma_s^2 \end{bmatrix}. \quad (2.3)$$

We let the pairwise correlation coefficient, ρ_s , vary between 0 and 1. When $T = 1$, one of the N possible location is chosen with equal probability and the stimuli at that location is assigned the target orientation. If $\mathbf{1}_j$ represents the N -dimensional vector having j^{th} entry as 1 and rest zeros, then

$$p(\mathbf{T}|T = 1) = \frac{1}{N} \sum_{j=1}^N \delta(\mathbf{T} - \mathbf{1}_j). \quad (2.4)$$

When the target is present at location j , for some fixed $j \in \{1, 2, \dots, N\}$, we have

$$p(s_j|T_j = 1) = \delta(s_j - s_T). \quad (2.5)$$

In such a case, the distractors will be present at all locations but the j^{th} one. Therefore, we can decompose the Eq. (2.4) into target and distractors' locations:

$$p(T_j = 1|T = 1) = \frac{1}{N} \delta(T_j - 1), \text{ and } p(\mathbf{T}_{\setminus j} = \mathbf{0}_{N-1}|T = 1) = \delta(\mathbf{T}_{\setminus j} - \mathbf{0}_{N-1}).$$

2.1. GENERATIVE MODEL

The vector $\mathbf{T}_{\setminus j}$ is obtained by removing the j^{th} location from the spatial location vector, \mathbf{T} . In the target present trial, the orientations of the $N - 1$ distractors are denoted by $\mathbf{s}_{\setminus j}$ (see notation in Appendix A). They are drawn from an $(N - 1)$ -dimensional multivariate normal distribution with $(N - 1)$ -dimensional mean vector $\mathbf{s}_{\mathbf{D}_{\setminus j}}$ and covariance, $\Sigma_{\mathbf{s}_{\setminus j}}$ (conditioned on j^{th} stimuli being the target). We write

$$\mathbf{s}_{\setminus j} | T = 1 \sim \mathcal{N}(\mathbf{s}_{\mathbf{D}_{\setminus j}}, \Sigma_{\mathbf{s}_{\setminus j}}), \quad (2.6)$$

where the $(N - 1) \times (N - 1)$ covariance matrix, $\Sigma_{\mathbf{s}_{\setminus j}}$ is obtained by removing the j^{th} row and the j^{th} column of $\Sigma_{\mathbf{s}}$.

We assume that an observer makes a measurement, x_i , of the presented stimulus, s_i , for $i \in \{1, 2, \dots, N\}$. It is commonly assumed that these measurements are noisy but unbiased and are normally distributed [94, 98, 99, 151]. Therefore, we assume that at each location i ,

$$x_i | s_i \sim \mathcal{N}(s_i, \sigma_i^2). \quad (2.7)$$

Moreover, we consider the measurement noise to be independent between locations. Hence, for the vector of measurements, $\mathbf{x} = (x_1, x_2, \dots, x_N)$, we write

$$\mathbf{x} | \mathbf{s} \sim \mathcal{N}(\mathbf{s}, \Sigma_{\mathbf{x}}) = \prod_{i=1}^N \mathcal{N}(s_i, \sigma_i^2). \quad (2.8)$$

Here $\Sigma_{\mathbf{x}}$ is an $N \times N$ diagonal matrix with entries $\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2$ on the diagonal. The optimal-observer model of the task is illustrated in Figure 2.1.




Figure 2.1: **Statistical structure of relevant task variables in the optimal-observer model for a target detection task with stimulus correlations. (A) Generative model.** The nodes represent the variables in the task and arrows indicate conditional dependencies between them. The binary variable, T represents target presence for $T = 1$ and absent when $T = 0$. The standard deviation, σ_s and the pairwise correlation coefficient, ρ_s determine the structure of the stimulus, $\mathbf{s} = (s_1, s_2, \dots, s_N)$ in the task. An observer makes a measurement, x_i , of each presented stimulus s_i . These measurements are assumed to be noisy and independent between locations. **(B) Inference process.** The optimal observer infers T by “inverting” the generative model. The observer computes a decision variable, $d(\mathbf{x})$ based on the measurements, \mathbf{x} and it is given by the log-posterior ratio between the two possibilities, $\log(p(T = 1|\mathbf{x})/p(T = 0|\mathbf{x}))$. The sign of $d(\mathbf{x})$ gives the optimal estimate of T and it is denoted by \hat{T} .

2.2 Inference process

The observer infers target presence based on the stimulus measurements, \mathbf{x} and knowledge of the process that generated the stimulus, also called the generative model (Figure 2.1(A)). Specifically, an optimal observer computes the probability of $T = 0$ and the probability of $T = 1$, given \mathbf{x} . The observer then chooses the option with highest probability. This is equivalent to computing the log posterior ratio,

$$d(\mathbf{x}) = \underbrace{\log \frac{p(T = 1|\mathbf{x})}{p(T = 0|\mathbf{x})}}_{\text{log posterior ratio}} \quad (2.9)$$

and reporting "target present" when $d(\mathbf{x}) > 0$, and "target absent" otherwise. The relation, $d(\mathbf{x}) > 0$ to make an estimate is known as the *Bayesian decision rule* and $d(\mathbf{x})$ itself is referred as the *Bayesian decision variable*. Here $d(\mathbf{x}) = 0$ represents the *decision boundary* and 0 is also called as the *decision criterion*.

By applying Bayes' theorem in the above equation we obtain

$$d(\mathbf{x}) = \underbrace{\log \frac{p(T = 1|\mathbf{x})}{p(T = 0|\mathbf{x})}}_{\text{log posterior ratio}} = \underbrace{\log \frac{p(\mathbf{x}|T = 1)}{p(\mathbf{x}|T = 0)}}_{\text{log likelihood ratio}} + \underbrace{\log \frac{p(T = 1)}{p(T = 0)}}_{\text{log prior ratio}}. \quad (2.10)$$

Here $p(T = 1)$ denotes the observer's prior belief that the target is present. Based on the above equation, the observer reports "target present" when the log-likelihood ratio is greater than the negative log prior ratio, *i.e.*,

$$\log \frac{p(\mathbf{x}|T = 1)}{p(\mathbf{x}|T = 0)} > -\log \frac{p(T = 1)}{p(T = 0)}.$$

Also, it is easy to see that any change in prior results in shifting of the decision criterion, thus prior has a large effect in the inference process of the observer. The

optimal observer uses a uniform prior on T to compute the log-likelihood ratio and the required decision variable. We denote the log-likelihood ratio for the task by $L_{ST}(\mathbf{x})$ and the decision variable by $d_{ST}(\mathbf{x})$.

2.2.1 The log-likelihood ratio

The observer needs to marginalize over intermediate variables, \mathbf{T} and \mathbf{s} to compute the log-likelihood ratio. The marginalization process is described in Section 1.3.4.1 and illustrated with an example in Section 1.4.1. In this case, we compute the log-likelihood ratio in the following manner:

$$\begin{aligned}
 L_{ST}(\mathbf{x}) &= \log \frac{p(\mathbf{x}|T=1)}{p(\mathbf{x}|T=0)} = \log \frac{\sum_{\mathbf{T}} p(\mathbf{x}|\mathbf{T}, T=1)p(\mathbf{T}|T=1)}{\sum_{\mathbf{T}} p(\mathbf{x}|\mathbf{T}, T=0)p(\mathbf{T}|T=0)} \\
 &= \log \frac{\sum_{\mathbf{T}} \left(p(\mathbf{x}|\mathbf{T}, T=1) \frac{1}{N} \sum_{j=1}^N \delta(\mathbf{T} - \mathbf{1}_j) \right)}{\sum_{\mathbf{T}} p(\mathbf{x}|\mathbf{T}, T=0)\delta(\mathbf{T} - \mathbf{0}_N)} \\
 &= \log \frac{1}{N} \frac{\sum_{j=1}^N \left(\sum_{\mathbf{T}} p(\mathbf{x}|\mathbf{T}, T=1)\delta(\mathbf{T} - \mathbf{1}_j) \right)}{\sum_{\mathbf{T}} p(\mathbf{x}|\mathbf{T}, T=0)\delta(\mathbf{T} - \mathbf{0}_N)} \\
 &= \log \frac{1}{N} \sum_{j=1}^N \frac{p(\mathbf{x}|\mathbf{T} = \mathbf{1}_j)}{p(\mathbf{x}|\mathbf{T} = \mathbf{0}_N)} \\
 &= \log \frac{1}{N} \sum_{j=1}^N \frac{\int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|\mathbf{T} = \mathbf{1}_j)d\mathbf{s}}{\int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|\mathbf{T} = \mathbf{0}_N)d\mathbf{s}}. \tag{2.11}
 \end{aligned}$$

2.2. INFERENCE PROCESS

Here δ is the generalized Kronecker delta function defined on \mathbb{R}^N :

$$\delta(\mathbf{z}) = \begin{cases} 1, & \text{if and only if } z_i = 0 \text{ for all } 1 \leq i \leq N, \\ 0, & \text{otherwise.} \end{cases}$$

We further simplify the above expressions by decomposing the vector \mathbf{s} into target stimulus, s_j and distractors, $\mathbf{s}_{\setminus j}$. Similarly, we break the vector \mathbf{x} into a target measurement, x_j and distractors measurements, $\mathbf{x}_{\setminus j}$. Similar to $\Sigma_{\mathbf{s}_{\setminus j}}$, we define matrix $\Sigma_{\mathbf{x}_{\setminus j}}$ obtained by removing the j^{th} row and column of matrix $\Sigma_{\mathbf{x}}$, so that

$$\mathbf{x}_{\setminus j} | \mathbf{s}_{\setminus j} \sim \mathcal{N}(\mathbf{s}_{\setminus j}, \Sigma_{\mathbf{x}_{\setminus j}}). \quad (2.12)$$

Thus, we obtain

$$\begin{aligned} L_{\text{ST}}(\mathbf{x}) &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{\int p(x_j | s_j) p(s_j | T_j = 1) p(\mathbf{x}_{\setminus j} | \mathbf{s}_{\setminus j}) p(\mathbf{s}_{\setminus j} | \mathbf{T}_{\setminus j} = \mathbf{0}_{N-1}) ds_j d\mathbf{s}_{\setminus j}}{\int p(\mathbf{x} | \mathbf{s}) p(\mathbf{s} | T = 0) d\mathbf{s}} \right] \\ &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{\int f(x_j; s_j, \sigma_j^2) \delta(s_j - s_T) f(\mathbf{x}_{\setminus j}; \mathbf{s}_{\setminus j}, \Sigma_{\mathbf{x}_{\setminus j}}) f(\mathbf{s}_{\setminus j}; \mathbf{s}_{\mathbf{D}_{\setminus j}}, \Sigma_{\mathbf{s}_{\setminus j}}) ds_j d\mathbf{s}_{\setminus j}}{\int f(\mathbf{x}; \mathbf{s}, \Sigma_{\mathbf{x}}) f(\mathbf{s}; \mathbf{s}_{\mathbf{D}}, \Sigma_{\mathbf{s}}) d\mathbf{s}} \right]. \end{aligned}$$

We now apply the product and integral rules for normal distributions in Eqs. (B.3) and (B.4), and denote

$$\mathbf{C} = \Sigma_{\mathbf{s}} + \Sigma_{\mathbf{x}} \text{ and } \mathbf{C}_{\setminus j} = \Sigma_{\mathbf{s}_{\setminus j}} + \Sigma_{\mathbf{x}_{\setminus j}}. \quad (2.13)$$

The matrix $\mathbf{C}_{\setminus j}$ can also be obtained by removing the j^{th} row and column of matrix \mathbf{C} . In the case of positive definite covariance matrices, *i.e.*, for $\rho_s \neq 1$, we integrate

2.2. INFERENCE PROCESS

and obtain the following expression for the log-likelihood ratio,

$$\begin{aligned}
 L_{\text{ST}}(\mathbf{x}) &= \log \frac{1}{N} \sum_{j=1}^N \frac{f(x_j; s_T, \sigma_j^2) f(\mathbf{x}_{\setminus j}; \mathbf{s}_{\mathbf{D}_{\setminus j}}, \mathbf{C}_{\setminus j})}{f(\mathbf{x}; \mathbf{s}_{\mathbf{D}}; \mathbf{C})} \\
 &= \log \left[\frac{1}{N} \sum_{j=1}^N \sqrt{\frac{|\mathbf{C}|}{\sigma_j^2 |\mathbf{C}_{\setminus j}|}} \exp \left(-\frac{(x_j - s_T)^2}{2\sigma_j^2} - \frac{1}{2} (\mathbf{x}_{\setminus j} - \mathbf{s}_{\mathbf{D}_{\setminus j}})^T \mathbf{C}_{\setminus j}^{-1} (\mathbf{x}_{\setminus j} - \mathbf{s}_{\mathbf{D}_{\setminus j}}) \right. \right. \\
 &\quad \left. \left. + \frac{1}{2} (\mathbf{x} - \mathbf{s}_{\mathbf{D}})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{s}_{\mathbf{D}}) \right) \right]. \tag{2.14}
 \end{aligned}$$

We further reduce the above expression by computing the determinant and inverse of matrices \mathbf{C} and $\mathbf{C}_{\setminus j}$. The matrices \mathbf{C} and $\mathbf{C}_{\setminus j}$ have rank 1, therefore, we use the matrix determinant lemma and the Sherman-Morrison formula to compute their determinants and inverses, respectively.

2.2.1.1 Determinants and inverses of matrices \mathbf{C} and $\mathbf{C}_{\setminus j}$

We decompose the matrix $\mathbf{C} = \begin{bmatrix} \sigma_s^2 + \sigma_1^2 & \rho_s \sigma_s^2 & \cdots & \rho_s \sigma_s^2 \\ \rho_s \sigma_s^2 & \sigma_s^2 + \sigma_2^2 & \cdots & \rho_s \sigma_s^2 \\ \vdots & & \ddots & \vdots \\ \rho_s \sigma_s^2 & \rho_s \sigma_s^2 & \cdots & \sigma_s^2 + \sigma_N^2 \end{bmatrix}$ as

$$\underbrace{\begin{bmatrix} \sigma_s^2(1 - \rho_s) + \sigma_1^2 & 0 & \cdots & 0 \\ 0 & \sigma_s^2(1 - \rho_s) + \sigma_2^2 & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_s^2(1 - \rho_s) + \sigma_N^2 \end{bmatrix}}_D + \rho_s \sigma_s^2 \underbrace{\begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix}}_J$$

2.2. INFERENCE PROCESS

and compute the determinant of \mathbf{C} using the matrix determinant lemma in Appendix B.2,

$$|\mathbf{C}| = |D + \rho_s \sigma_s^2 J| = |D + \rho_s \sigma_s^2 \mathbf{1} \mathbf{1}^\top| = (1 + \mathbf{1}^\top D^{-1} \mathbf{1}) |D|,$$

where $\mathbf{1}$ is a column vector of ones.

Since D is a diagonal matrix, $|D| = \prod_{i=1}^N (\sigma_s^2(1 - \rho_s) + \sigma_i^2)$ and

$$(D^{-1})_{k,l} = \begin{cases} \frac{1}{\sigma_s^2(1 - \rho_s) + \sigma_k^2}, & \text{if } k = l, \\ 0, & \text{otherwise.} \end{cases}$$

We also define

$$w_i = \frac{1}{\sigma_i^2}, \quad \tilde{w}_i = \frac{1}{\sigma_s^2(1 - \rho_s) + \sigma_i^2}, \quad \tilde{W} = \sum_{i=1}^N \tilde{w}_i, \quad \text{and } \tilde{W}_{\setminus j} = \sum_{i \neq j}^N \tilde{w}_i, \quad (2.15)$$

and obtain

$$|\mathbf{C}| = \left(1 + \rho_s \sigma_s^2 \tilde{W}\right) \prod_{i=1}^N \frac{1}{\tilde{w}_i}. \quad (2.16)$$

Similarly, we compute

$$|\mathbf{C}_{\setminus j}| = \left(1 + \rho_s \sigma_s^2 \tilde{W}_{\setminus j}\right) \prod_{i \neq j}^N \frac{1}{\tilde{w}_i}. \quad (2.17)$$

Next, we compute the inverse of matrix \mathbf{C} using the Sherman-Morrison Formula described in Appendix B.2. Specifically, we obtain

$$\mathbf{C}^{-1} = D^{-1} - \frac{D^{-1} \mathbf{1} \mathbf{1}^\top D^{-1}}{1 + \mathbf{1}^\top D^{-1} \mathbf{1}} = \begin{bmatrix} \tilde{w}_1 - \alpha \tilde{w}_1^2 & -\alpha \tilde{w}_1 \tilde{w}_2 & \cdots & -\alpha \tilde{w}_1 \tilde{w}_N \\ -\alpha \tilde{w}_2 \tilde{w}_1 & \tilde{w}_2 - \alpha \tilde{w}_2^2 & \cdots & -\alpha \tilde{w}_2 \tilde{w}_N \\ \vdots & \vdots & \ddots & \vdots \\ -\alpha \tilde{w}_N \tilde{w}_1 & -\alpha \tilde{w}_N \tilde{w}_2 & \cdots & \tilde{w}_N - \alpha \tilde{w}_N^2 \end{bmatrix}, \quad (2.18)$$

2.2. INFERENCE PROCESS

where

$$\alpha = \frac{\rho_s \sigma_s^2}{1 + \rho_s \sigma_s^2 \tilde{W}} = \frac{1}{\frac{1}{\rho_s \sigma_s^2} + \tilde{W}}. \quad (2.19)$$

The inverse of matrix $\mathbf{C}_{\setminus j}$ has the same form as \mathbf{C}^{-1} in Eq. (2.18) except that α is replaced by $\alpha_{\setminus j}$ with

$$\alpha_{\setminus j} = \frac{\rho_s \sigma_s^2}{1 + \rho_s \sigma_s^2 \tilde{W}_{\setminus j}} = \frac{1}{\frac{1}{\rho_s \sigma_s^2} + \tilde{W}_{\setminus j}}. \quad (2.20)$$

2.2.1.2 Bayesian decision variable

We continue simplifying the log-likelihood ratio (Eq. (2.14)) in order to obtain an analytically tractable expression for the decision variable, $d_{\text{ST}}(\mathbf{x})$ (Eq. (2.9)). First, we compute the required ratio in Eq. (2.14) of determinants of the two matrices

$$\frac{|\mathbf{C}|}{\sigma_j^2 |\mathbf{C}_{\setminus j}|} = \frac{w_j |\mathbf{C}|}{|\mathbf{C}_{\setminus j}|} = \frac{w_j (1 + \rho_s \sigma_s^2 \tilde{W}) \prod_i^N \frac{1}{\tilde{w}_i}}{(1 + \rho_s \sigma_s^2 \tilde{W}_{\setminus j}) \prod_{i \neq j}^N \frac{1}{\tilde{w}_i}} = \frac{w_j (1 + \rho_s \sigma_s^2 \tilde{W})}{\tilde{w}_j (1 + \rho_s \sigma_s^2 \tilde{W}_{\setminus j})} = \frac{w_j \alpha_{\setminus j}}{\tilde{w}_j \alpha}.$$

Next, we compute the exponent terms in Eq. (2.14):

$$\begin{aligned} (\mathbf{x} - \mathbf{s}_D)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{s}_D) &= \sum_{i=1}^N (\tilde{w}_i - \alpha \tilde{w}_i^2) (x_i - s_D)^2 - \alpha \sum_{i \neq k}^N \tilde{w}_i \tilde{w}_k (x_i - s_D) (x_k - s_D), \\ (\mathbf{x}_{\setminus j} - \mathbf{s}_{D_{\setminus j}})^T \mathbf{C}_{\setminus j}^{-1} (\mathbf{x}_{\setminus j} - \mathbf{s}_{D_{\setminus j}}) &= \sum_{i \neq j}^N (\tilde{w}_i - \alpha_{\setminus j} \tilde{w}_i^2) (x_i - s_D)^2 \\ &\quad - \alpha_{\setminus j} \sum_{\substack{i \neq k \\ i, k \neq j}}^N \tilde{w}_i \tilde{w}_k (x_i - s_D) (x_k - s_D), \end{aligned}$$

and combine them to obtain

$$\begin{aligned}
 & - (\mathbf{x}_{\setminus j} - \mathbf{s}_{D_{\setminus j}})^T \mathbf{C}_{\setminus j}^{-1} (\mathbf{x}_{\setminus j} - \mathbf{s}_{D_{\setminus j}}) + (\mathbf{x} - \mathbf{s}_D)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{s}_D) \\
 & = (\tilde{w}_j - \alpha \tilde{w}_j^2) (x_j - s_T)^2 - 2\tilde{w}_j \alpha (x_j - s_D) \sum_{i \neq j}^N \tilde{w}_i (x_i - s_D) \\
 & + (\alpha_{\setminus j} - \alpha) \sum_{i, k \neq j}^N \tilde{w}_i \tilde{w}_k (x_i - s_D) (x_k - s_D) \\
 & = \frac{\tilde{w}_j \alpha}{\alpha_{\setminus j}} (x_j - s_T)^2 - 2\tilde{w}_j \alpha (x_j - s_D) \sum_{i \neq j}^N \tilde{w}_i (x_i - s_D) + \tilde{w}_j \alpha \alpha_{\setminus j} \left(\sum_{i \neq j}^N \tilde{w}_i (x_i - s_D) \right)^2.
 \end{aligned}$$

We substitute the above expressions in Eq. (2.14) to compute the log-likelihood ratio and therefore, obtain the following expression for the decision variable

$$\begin{aligned}
 d_{ST}(\mathbf{x}) = \log & \left[\frac{1}{N} \sum_{j=1}^N \sqrt{\frac{w_j (1 + \rho_s \sigma_s^2 \tilde{W})}{\tilde{w}_j (1 + \rho_s \sigma_s^2 \tilde{W}_{\setminus j})}} \exp \left(-\frac{1}{2} \underbrace{w_j (x_j - s_T)^2}_I \right) \right. \\
 & + \frac{1}{2} \underbrace{(\tilde{w}_j - \alpha \tilde{w}_j^2) (x_j - s_D)^2}_{II} - \underbrace{\tilde{w}_j \alpha (x_j - s_D) \sum_{i \neq j}^N \tilde{w}_i (x_i - s_D)}_{III} \\
 & \left. + \frac{1}{2} \underbrace{(\alpha_{\setminus j} - \alpha) \sum_{i, k \neq j}^N \tilde{w}_i \tilde{w}_k (x_i - s_D) (x_k - s_D)}_{IV} \right]. \tag{2.21}
 \end{aligned}$$

The above equation gives the non-linear decision variable $d_{ST}(\mathbf{x})$ in terms of the stimulus measurement and model parameters: the total number of stimuli, N , the variance, and covariance between distractors orientations given by σ_s^2 and ρ_s . These parameters govern the statistical structure of the visual stimuli. The observer must infer these model parameters in order to make an optimal decision on the task. We assume that the observer is aware of the noise with which the measurement is made.

2.2.2 Interpretation of the decision variable

We note that the decision variable computed in Eq. (2.21) depends in an intricate manner on the model parameters that describe the structure of the stimulus and its measurements. Although the expression is complex, each term in the exponent has an intuitive interpretation. We can think of the different terms as different pieces of evidence about whether the j^{th} stimulus is a target: (I) if the j^{th} measurement is close to the target orientation, this term is larger (less negative), corresponding to an increased likelihood that the j^{th} stimulus is the target; (II) the second term decreases as the j^{th} measurement approaches the mean distractor orientation, this corresponds to a decreased likelihood that the j^{th} stimulus is the target; (III) the third term compares the j^{th} measurement to the sample distractor mean; if it is large, it is less likely that the j^{th} stimulus is the target; and (IV) the fourth term can be rewritten in terms of sample covariance of potential distractor measurements, and in that case a large covariance increases the likelihood that the j^{th} stimulus is the target. Therefore, different terms in Eq. (2.21) contribute towards finding the target.

2.3 MAP estimate distribution

We denote the observer's MAP estimate of T by \hat{T} . The optimal observer responds \hat{T} based on the sign of the decision variable $d_{\text{ST}}(\mathbf{x})$ computed in Eq. (2.21). The probability of the optimal observer responding \hat{T} given a fixed stimuli $\mathbf{s}_{\text{fixed}}$ is

denoted by $p(\hat{T}|\mathbf{s}_{\text{fixed}})$. As illustrated in [151], we compute this probability by marginalizing over the hypothesized observations \mathbf{x} generated by \mathbf{s} :

$$p(\hat{T}|\mathbf{s}_{\text{fixed}}) = \int p(\hat{T}|\mathbf{x})p(\mathbf{x}|\mathbf{s}_{\text{fixed}})d\mathbf{x} = \int \delta_{\hat{T},\text{sgn}(d_{\text{ST}}(\mathbf{x}))} p(\mathbf{x}|\mathbf{s}_{\text{fixed}})d\mathbf{x}. \quad (2.22)$$

This means that the probability of reporting $\hat{T} = 1$ is evaluated by averaging over the observations that are drawn from the distribution $\mathbf{x}|\mathbf{s}_{\text{fixed}}$ and have $d_{\text{ST}}(x) > 0$. But, we note that the decision variable, $d_{\text{ST}}(\mathbf{x})$, is a non-linear function of \mathbf{x} and hence, the above expression is analytically intractable. Therefore, we use numerical approximations for our computational purposes. In particular, we apply *Monte Carlo* method that is described in Section 3.2.2.

2.4 Suboptimal models

The computations involved in evaluating the exponent terms in Eq. (2.21) are complex and require complete information about the generative model. But in general, an observer may not learn and use the correct generative model in making decisions. The observer can either use an incorrect assumption about the correlation strength, ρ_s or may not use an equal odd prior for T . In that case, their inferences will be suboptimal. Therefore, we need to investigate what inference models are used by the observers to make their decisions and what model parameters have been assumed by them to infer their responses. We study a range of suboptimal models in Chapter 4. For mathematical purposes, we consider two special variants of the suboptimal models below. These are the cases of homogeneous and heterogeneous distractors. The experimental studies [94, 99] have

well characterized the human behavior on both these conditions; however, they assumed different probability distributions in the generative model.

2.4.1 Heterogeneous model, $\rho_s = 0$

In this model, we assume that the observer does not learn any information about the stimulus correlations and makes decisions assuming that no structure is present in the scenes. This amounts to the observer using $\rho_s = 0$ in making decisions; therefore, the decision variable in Eq. (2.21) reduces to the following simplified expression:

$$\begin{aligned} d_{ST}(\mathbf{x}) &= \log \left[\frac{1}{N} \sum_{j=1}^N \sqrt{(1 + w_j \sigma_s^2)} \exp \left(-\frac{1}{2} w_j (x_j - s_T)^2 + \frac{1}{2} \tilde{w}_j (x_j - s_D)^2 \right) \right] \\ &= \log \left[\frac{1}{N} \sum_{j=1}^N \sqrt{\frac{\sigma_j^2 + \sigma_s^2}{\sigma_j^2}} \exp \left((s_T - s_D) \frac{x_j - \frac{(s_T + s_D)}{2}}{\sigma_j^2 + \sigma_s^2} - \frac{\sigma_s^2 (x_j - s_T)^2}{2\sigma_j^2 (\sigma_j^2 + \sigma_s^2)} \right) \right]. \end{aligned} \quad (2.23)$$

This condition reflects a suboptimal decision and we would expect the observer following this model to have a low performance.

2.4.2 Homogeneous model, $\rho_s = 1$

We also consider another extreme possibility: the observer may assume that the stimuli are always maximally correlated and may make decisions using $\rho_s = 1$. That is, the distractors always have a common orientation. We note that when $\rho_s = 1$, the covariance matrix Σ_s (defined in Eq. (2.3)) is singular. Therefore, the

2.4. SUBOPTIMAL MODELS

conditions to use Eqs. (B.3) and (B.4) are violated. Hence, we cannot use Eq. (2.21) here.

Instead, we independently compute the decision variable in this case. We note that the covariance matrix, Σ_s reduces to a single term σ_s^2 for $\rho_s = 1$, and the multivariate normal distribution to a one-dimensional Gaussian distribution with mean, s_D and variance, σ_s^2 . This results in having all distractors with an identical orientation and we denote that common orientation by s . Thus, the common distractor orientation follows the Gaussian distribution

$$s|T \sim \mathcal{N}(s_D, \sigma_s^2).$$

As before, we compute the log-likelihood ratio by marginalizing over the intermediate variables \mathbf{T} and \mathbf{s} as in Eq. (2.11) to obtain

$$L_{\text{ST}}(\mathbf{x}) = \log \frac{p(\mathbf{x}|T=1)}{p(\mathbf{x}|T=0)} = \log \frac{1}{N} \sum_{j=1}^N \frac{\int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|\mathbf{T}=\mathbf{1}_j)d\mathbf{s}}{\int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|\mathbf{T}=\mathbf{0}_N)d\mathbf{s}}.$$

We now use the assumption about the independence (between locations) of measurement noise and the product form in Eq. (2.8) for the probability distribution $P(\mathbf{x}|\mathbf{s})$. Also, we decompose the measurement vector \mathbf{x} into a target measurement, x_j and a common distractor measurement, x_i for $i \neq j$. This gives us

2.4. SUBOPTIMAL MODELS

$$\begin{aligned}
 L_{\text{ST}}(\mathbf{x}) &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{\int p(x_j|s_j = s_T) p(s_j|T_j = 1) ds_j \left(\prod_{i \neq j}^N p(x_i|s_i = s) \right) p(s|T) ds}{\int \left(\prod_{i=1}^N p(x_i|s_i = s) \right) p(s|T) ds} \right] \\
 &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{\int f(x_j; s_j, \sigma_j^2) \delta(s_j - s_T) ds_j \left(\prod_{i \neq j}^N f(x_i; s, \sigma_i^2) \right) f(s; s_D, \sigma_s^2) ds}{\int \left(\prod_{i=1}^N f(x_i, s, \sigma_i^2) \right) f(s; s_D, \sigma_s^2) ds} \right].
 \end{aligned}$$

We use Eqs. (B.2) and (B.4) to further compute the expression for the log-likelihood ratio

$$\begin{aligned}
 L_{\text{ST}}(\mathbf{x}) &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{f(x_j; s_T, \sigma_j^2)}{\sqrt{(2\pi)^{(N-1)} \left(\frac{1}{\sigma_s^2} + \sum_{i \neq j}^N \frac{1}{\sigma_i^2} \right) \sigma_s \prod_{i \neq j}^N \sigma_i}} \exp \left(-\frac{1}{2} \left(\sum_{i \neq j}^N \frac{x_i^2}{\sigma_i^2} + \frac{s_D^2}{\sigma_s^2} - \frac{\left(\sum_{i \neq j}^N \frac{x_i}{\sigma_i^2} + \frac{s_D}{\sigma_s^2} \right)^2}{\frac{1}{\sigma_s^2} + \sum_{i \neq j}^N \frac{1}{\sigma_i^2}} \right) \right)}{\sqrt{(2\pi)^N \left(\frac{1}{\sigma_s^2} + \sum_{i=1}^N \frac{1}{\sigma_i^2} \right) \sigma_s \prod_{i=1}^N \sigma_i}} \exp \left(-\frac{1}{2} \left(\sum_{i=1}^N \frac{x_i^2}{\sigma_i^2} + \frac{s_D^2}{\sigma_s^2} - \frac{\left(\sum_{i=1}^N \frac{x_i}{\sigma_i^2} + \frac{s_D}{\sigma_s^2} \right)^2}{\frac{1}{\sigma_s^2} + \sum_{i=1}^N \frac{1}{\sigma_i^2}} \right) \right)} \right].
 \end{aligned}$$

2.4. SUBOPTIMAL MODELS

We simplify above expression to obtain

$$\begin{aligned}
 L_{ST}(\mathbf{x}) = \log & \left[\frac{1}{N} \sum_{j=1}^N \sqrt{\frac{\left(\frac{1}{\sigma_s^2} + \sum_{i=1}^N \frac{1}{\sigma_i^2} \right)}{\left(\frac{1}{\sigma_s^2} + \sum_{i \neq j}^N \frac{1}{\sigma_i^2} \right)}} \right. \\
 & \times \exp \left(-\frac{(x_j - s_T)^2}{2\sigma_j^2} + \frac{x_j^2}{2\sigma_j^2} - \frac{1}{2} \frac{\left(\sum_{i=1}^N \frac{x_i}{\sigma_i^2} + \frac{s_D}{\sigma_s^2} \right)^2}{\frac{1}{\sigma_s^2} + \sum_{i=1}^N \frac{1}{\sigma_i^2}} + \frac{1}{2} \frac{\left(\sum_{i \neq j}^N \frac{x_i}{\sigma_i^2} + \frac{s_D}{\sigma_s^2} \right)^2}{\frac{1}{\sigma_s^2} + \sum_{i \neq j}^N \frac{1}{\sigma_i^2}} \right) \left. \right]. \quad (2.24)
 \end{aligned}$$

Assuming the observer uses a uniform prior over T , the above equation represents the decision variable under the assumption of heterogeneous distractors. Using the variables defined in Eqs. (2.15) and (2.19), we rewrite the following compact expression for the decision variable in the case of $\rho_s = 1$:

$$\begin{aligned}
 d_{ST}(\mathbf{x}) = \log & \left[\frac{1}{N} \sum_{j=1}^N \sqrt{(1 + w_j \alpha_{\setminus j})} \exp \left(-\frac{1}{2} \left(w_j (x_j - s_T)^2 - w_j x_j^2 \right. \right. \right. \\
 & \left. \left. \left. + \alpha \left(\sum_{i=1}^N w_i x_i + \frac{s_D}{\sigma_s^2} \right)^2 - \alpha_{\setminus j} \left(\sum_{i \neq j}^N w_i x_i + \frac{s_D}{\sigma_s^2} \right)^2 \right) \right) \right]. \quad (2.25)
 \end{aligned}$$

The above decision variable characterizes the decision-making behavior of an observer that assumes the distractors to be always perfectly correlated. The observer making decisions based on this variable will be suboptimal since the correct correlation strength, ρ_s is ignored and always assumed as 1.

The homogeneous and heterogeneous models explained above are two particular suboptimal models. We consider a range of other possible suboptimal models

in Chapter 4. We now provide the details of the experiment designed based on the target detection task.

2.5 Experimental Methods

We conducted an experiment based on the design of the target detection task described in Section 2.1. The experiment was performed under the supervision of Dr. Wei Ji Ma in the Theoretical Systems Neuroscience laboratory at the Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA.

The aim of our experimental study was to determine whether human observers use the structures present in the visual scenes to infer their decisions. In the case that they do, we want to examine whether they are able to infer the correct correlation strength, ρ_s that is used to generate the experimental displays. If not, we test several suboptimal models in Chapter 4 and find the one that best explain the responses of the subjects on the experiment. In the following section, we provide the details of the experiment.

2.5.1 Subjects

Eleven subjects (6 males and 5 females) participated in the experiment. All subjects had normal or corrected-to-normal acuity and gave informed consent.

2.5.2 Apparatus and stimuli

Stimuli were generated in Matlab using the Psychophysics Toolbox [24, 113] and were presented on a 21" LCD monitor with a refresh rate of 60 Hz. Subjects viewed the displays from a distance of approximately 60 cm. The background luminance was 33.1 cd/m^2 . A set of 4 stimuli were shown on each trial. On target-present trials, the stimulus set consisted of one target and 3 distractors while on target-absent trials, it contained 4 distractors. A target was present in exactly half of the trials. Each stimulus was a Gabor patch (or a Gabor filter) with a spatial frequency of approximately 2.67 cycles/deg, a standard deviation of 0.26 deg, and a peak luminance of 136 cd/m^2 . A Gabor patch is a sine wave multiplied by a Gaussian function (additional details can be found in Appendix C). Figure 2.2(A) shows one example of a Gabor patch. Stimuli were placed on a circle centered at the fixation cross with a radius of 3.2 degrees of visual angle. The position of the first stimulus was chosen at random on each trial and other stimuli were placed in a way so that the angular distance between two adjacent stimuli was always 45° . The target (s_T) and mean distractor orientation (s_D) were set to vertical and used to define the origin. The standard deviation of the distractor distribution σ_s was fixed at 15° while the correlation coefficient ρ_s was varied to be $0, \frac{1}{3}, \frac{2}{3}$, and 1 across different experimental sessions.

2.5.3 Sessions and blocks

The experiment was split into four different sessions. The correlation coefficient was fixed at 0, 1/3, 2/3, or 1 within a session. The order of the sessions was randomized across subjects. Each session consisted of one practice block of 50 trials and 6 testing blocks of 150 trials each and lasted for about 50 minutes. A 30 sec break was provided between blocks. After each block, performance on that block was revealed to the subject along with the scores of the other subjects who had completed the same session. Each subject completed a total of 3600 test trials. All subjects were instructed about the experiment at the beginning of their first session with a demo consisting of 10 practice trials.

2.5.4 Procedure

Testing trials: Each test trial began with the display of a fixation cross at the center of the screen (0.5 sec), followed by the stimulus display containing 4 stimuli (0.1 sec), and followed by a screen with the fixation cross until the subject responded. The subject reported whether a target was present or absent through a key press. Feedback was provided by subsequently coloring the fixation cross green (correct) or red (incorrect) during the (0.75 sec) inter-trial period. The experiment and time course are shown in Figure 2.2(B).

Practice trials: Each practice trial was identical to a test trial, except that it was followed by a feedback screen showing the original stimulus display with a blue




Figure 2.2: **Target detection experiment procedure.** **(A) Gabor patch.** Subjects were presented with 4 stimuli on each trial. Each stimulus was a Gabor patch. The figure illustrates an example of a Gabor patch from the experiment. **(B) Time course of a test trial.** The experiment started with a display of fixation cross in the center followed by the stimulus display for 0.1 sec. Subjects reported through a key press whether a vertical stimulus was present in the display. After their response, a display screen was shown with a green (correct) or red (incorrect) fixation cross to provide feedback. **(C) Time course of a practice trial.** Each experimental session started with 50 practice trials. The procedure of a practice trial was same as a test trial, except that an additional feedback screen was shown for 2 sec at the end of the trial. The extra display contained the original stimulus with a blue circle marking the target stimulus if it was present. **(D) Sample displays from different experimental conditions.** The experiment was divided in four different sessions. Each session was characterized by the unique value of the correlation coefficient $\rho_s \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$ used in generating stimuli. The order of the sessions was randomized across subjects. This figure shows example displays from each experimental session.

2.5. EXPERIMENTAL METHODS

circle identifying the target stimulus, when present (Figure 2.2(C)). The data from practice trials were excluded for the analysis and results purposes.

Chapter 3

Model fitting and model comparison

Our perception about the true state of the world is based on our sensory information along with the prior knowledge of the world state. But it is not clear what computations the brain performs to combine the information it receives and makes a decision. Numerous theoretical and experimental studies provide evidence that our perception can be described as a process of probabilistic inference. In particular, Bayesian models are applied to study human behavior on various perceptual tasks. Several experimental studies have found that humans are near-optimal in simple visual perceptual tasks [68, 43, 8, 42, 99, 111, 7, 105, 94, 10, 83, 153, 98]. We note that observers need to have complete knowledge of the underlying generative model of the task in order to make the best possible decisions. However, we may expect that on a complex task or a task with a large number of latent variables, it may not always be possible for the observers to determine the correct generative model. Therefore, the relevant question here is what model of

the world observers use to make their decisions and what are the possible techniques that we can use to compute the predictions of a model for the experimental data.

In this chapter, we discuss some commonly used techniques of fitting a model to the responses of subjects on a psychophysical task. We describe the maximum-likelihood parameter estimation method to find the model parameters that describe the best fit to an experimental data. Further, we discuss the criteria that can be used to compare models and find the one that best describes the data. We present these techniques in a generalized form here, but discuss their applications, and possible issues in the context of our experimental study (described in Chapter 2). Specific to our experiment, we begin this chapter with a detailed description of the experimental data we have collected. Further, we elaborate the types of psychometric curves we use to represent the subjects' responses from our experiment. In the end, we describe how we use different model comparison criteria specific to the models used to predict our experimental data.

3.1 Experimental data

The set up of our experiment has been described in Section 2.5. Subjects were presented with a set of 4 stimuli on each experimental trial. The orientations of the stimuli were drawn randomly across trials. Also, the stimuli were placed at random locations. As a consequence, each subject was presented with a unique set of stimuli over the course of the experiment. We thus recorded the collection

3.1. EXPERIMENTAL DATA

of stimulus orientations presented to the subject. We then recorded whether any stimulus is a target in the trial, that is, we noted if $T = 1$ or $T = 0$. Furthermore, we documented the response of the subject in each trial as either 0 or 1. Finally, we measured the exact duration of the stimulus presentation on the display screen. We note that this duration is actually a constant (100 ms) in the experiment. In summary, for each subject we recorded the following information on each trial:

1. a set of 4 stimulus orientations, \mathbf{s} ,
2. the information about the target presence variable T ,
3. the subject's response or the MAP estimate \hat{T} ,
4. whether \hat{T} matches T or not, and
5. the duration of the stimulus presentation.

In the following section, we discuss the methods of obtaining model predictions for an experimental data and to use these predictions for fitting the subject's responses. In our analysis, we evaluate the predictions of a model for each individual instead of relying on average (over subjects) statistics. Therefore, our model fitting and model comparison processes are computed based on individual responses.

3.2 Model predictions

In the experiment, we record the responses of a subject - whether the subject responded "target present" or "absent" on each trial. However, we are not aware of the sensory measurements with which the subject made a decision on a trial. Thus, as part of the modeling process, we need to consider what measurements would have led the subject to make a particular response on the trial. Hence, we must use a reasonable assumption about the distribution of the measurements. Moreover, in general, we have no means of measuring what parameters the subject would have used to make the decision. For instance, the sensory noise with which the subject made the measurement is unknown to us and we need to estimate it in order to understand the behavior of the subject. Such parameters are sometimes referred as *free parameters* of the model and can be estimated from the data. Therefore, we consider two important issues here: to find the maximum-likelihood estimates of model parameters based on subject's data and to make predictions of the model for the data given those parameters.

Let us consider a model M having a parameter θ . The parameter θ can either be a scalar or vector quantity given the model. We assume that we have access to a subject's responses on K trials in the experiment. We denote the subject's response by a binary variable r_i and the set of presented stimuli by \mathbf{s}_i on the i^{th} trial. We are interested in evaluating the model prediction for the response, r_i on i^{th} trial given the model parameter θ . We denote $p(r_i|\mathbf{s}_i, M, \theta)$ as the probability of the subject response r_i given the stimuli on the i^{th} trial under the model M with hypothesized

parameter value θ . This is also known as the *response probability* under the model M with parameter θ .

3.2.1 Computing response probabilities

We would like to compute the response probability $p(r_i|\mathbf{s}_i, M, \theta)$ on the i^{th} trial under model M . This is similar to evaluating the distribution of the MAP estimate in Eq. (2.22) under the decision variable of the model M . The cases where such an evaluation is not analytically possible, we use Monte Carlo method [62, 11, 60, 72, 126, 15] for a numerical approximation of the involved integral. The integral is approximated by the sum which converges to the correct value as the number of measurement samples increases:

$$\begin{aligned} p(r_i|\mathbf{s}_i, M, \theta) &= \int p(r_i|\mathbf{x})p(\mathbf{x}_i|\mathbf{s}_i, M, \theta)d\mathbf{x} \approx \sum_{\mathbf{x}_i} p(r_i|\mathbf{x})p(\mathbf{x}_i|\mathbf{s}_i, M, \theta) \\ &= \sum_{\mathbf{x}_i} \delta_{r_i, \text{sgn}(d_M(\mathbf{x}_i))} p(\mathbf{x}_i|\mathbf{s}_i, M, \theta). \end{aligned} \quad (3.1)$$

Here \mathbf{x}_i is the hypothesized measurement of the subject on the i^{th} trial and $d_M(\mathbf{x})$ represents the Bayesian decision variable under the model M with parameter θ .

3.2.2 Monte Carlo algorithm

We now describe the Monte Carlo algorithm we implemented for computing the response probabilities under the assumptions of a model M .

1. The first step is to describe the model M and its parameter(s) θ .

3.2. MODEL PREDICTIONS

2. We then derive the Bayesian decision variable, $d_M(\mathbf{x})$ under the hypothesis of the model M .
3. We fix a value of the model parameter θ .
4. For this fixed hypothesized value of the parameter θ , we draw hypothesized measurements, \mathbf{x}_i of the subject following Eq. (2.8) on the i^{th} trial in response to the presented stimulus, \mathbf{s}_i . The stimuli \mathbf{s}_i are used from the experimental data, and not generated during this step.
5. Using the hypothesized measurements, \mathbf{x}_i , we evaluate the decision variable $d_M(\mathbf{x}_i)$ of the model M . Further, we compute the prediction for the subject's response, denoted by \hat{r}_i , on the i^{th} trial based on the decision rule for the model M .
6. We then compare the model predicted response \hat{r}_i with the subject's true response r_i on the i^{th} trial. A match between the two results in an increased probability of the response $p(r_i|\mathbf{s}_i, M, \theta)$.
7. We repeat steps 4 to 6 with R samples of measurements \mathbf{x}_i and thus compute R values of model predicted responses \hat{r}_i . We subsequently match each of them with the corresponding actual response r_i of the subject.
8. We obtain an approximation of the probability of response $p(r_i|\mathbf{s}_i, M, \theta)$ by averaging over the number of correct matches in R samples.
9. Next, we pick a different value of θ and evaluate the response probability on the i^{th} trial for a different parameter value by following steps 4 to 8.

Steps 3 to 9 are repeated for every experimental trial to obtain the likelihood function of the subject's actual responses under the model M with parameter θ .

3.3 Maximum-likelihood estimation

Our aim is to find the value of the model parameter θ that maximizes the response probabilities over all trials. In other words, we are interested in finding the model parameter at which the predictions of the model provide the best possible explanation for the behavior of subjects in an experiment. This amounts to finding the maximum-likelihood estimate of θ . Below, we discuss the parameter estimation method illustrated in [93, 94, 98]. The likelihood function of a parameter value θ is defined as the probability of the data given the model M with parameter θ :

$$L_M(\theta) = p(\text{data}|M, \theta).$$

For simplification purposes, it is generally assumed that the noise in the observer's responses is independent across trials. Thus, we can write the probability of the data given the model and its parameters as a product of probabilities over trials:

$$L_M(\theta) = \prod_{i=1}^K p(r_i|\mathbf{s}_i, M, \theta). \quad (3.2)$$

Here K denotes the total number of trials in the experiment. Maximizing the above parameter likelihood is equivalent to maximizing its logarithm $\log L_M(\theta)$,

$$\log L_M(\theta) = \sum_{i=1}^K \log p(r_i|\mathbf{s}_i, M, \theta). \quad (3.3)$$

The logarithm prevents possible numerical issues which could arise because of a very small probability that could mistakenly be treated as zero and could lead

3.3. MAXIMUM-LIKELIHOOD ESTIMATION

to potential errors in the computation of the product in Eq. (3.2). We denote the maximum-likelihood estimate (MLE) of parameter θ by $\hat{\theta}$ and the corresponding maximum value of the likelihood function as $L_M^* = \max_{\theta} L_M(\theta) = L_M(\hat{\theta})$.

In the event that an analytical expression is unavailable for the response probability $p(r_i|\mathbf{s}_i, M, \theta)$ under the model M and hence, for the log-likelihood function in Eq. (3.3), we evaluate the function using numerical methods. Therefore, most of our data analysis practices rely on obtaining an accurate approximation of the log-likelihood function. As expected, the accuracy of this evaluation depends on the size of the data and the number of measurement samples used in the numerical approximation. Further, we need to find an appropriate numerical algorithm to find the maximum of the numerically evaluated stochastic log-likelihood function $L_M(\theta)$. We use suitable optimization algorithms to find the maximum-likelihood estimate of the parameter θ that maximizes the log-likelihood function defined in Eq. (3.3).

In our target detection task (Chapter 2), we do not have any analytical approximation of the log-likelihood function for any model. Therefore, we numerically evaluate the function for a model M at all hypothesized values of the model parameter θ . But, the evaluation of our function, $\log L_M(\theta)$ depends on the decision variable, $d_{\text{ST}}(\mathbf{x})$ computed in Eq. (2.21), and thus it is a non-linear, non-smooth stochastic function. To maximize this stochastic function, we tested the following optimization algorithms and chose the one that was most suitable for our purposes.

3.4 Optimization techniques

In order to find the maximum of our non-linear stochastic objective function $L_M(\theta)$, we tried several different optimization algorithms and compared the obtained results. We checked the common practice of grid search method and the standard techniques such as *genetic algorithm* and *pattern search*. However, most of our models (described in Chapter 4) were high-dimensional in parameter space and thus optimization algorithms required tremendous amount of computational power and time to produce results. This restricted us to only use grid search method for our results purposes. In the following section, we briefly describe the three algorithms we had tested for our data.

3.4.1 Exhaustive or grid search

Exhaustive or brute-force is one of the simplest possible methods of optimizing an objective function by manually evaluating it on a predefined set of parameter space. The possible parameter space is identified and systematically divided into possibly a large set of discrete values known as the *grid space*. The objective function is then evaluated at all points of the pre-defined grid space and the point at which the function attains the global maximum is regarded as the maximum-likelihood estimate of the parameter. The accuracy of this method greatly depends on the choice of the parameter space and the grid spacing.

3.4. OPTIMIZATION TECHNIQUES

Most of our models (described in Chapter 4) have a high-dimensional parameter vector θ . This resulted in a multi-dimensional grid spacing for θ and also a large number of grid points. Also, the number of evaluations of the log-likelihood function has an exponential increase with respect to finer discretization of the grid space. Moreover, our likelihood function is evaluated using Monte Carlo method (see Eqs. (3.1) and (3.3)) and require a large number of measurement samples to guarantee convergence. As a consequence, finding maximum-likelihood parameter estimates for our models requires a large amount of computational resources and time. Furthermore, the precision of the results depends on the convergence of the log-likelihood evaluation and the discretization of the grid space.

However, we still use grid search method over other optimization algorithms to obtain our results (in Chapter 4). It is because this method can be run in parallel for each grid point as the evaluation of the log-likelihood function for our models is independent between grid points. Therefore, we can numerically evaluate the function at multiple grid points at the same time and combine the results together. To be able to run in parallel is the most crucial feature of this exhaustive search that led us to use this method for our analysis purposes. Also, we obtained consistent results with this approach. We performed several tests on synthetic data sets to ensure the accuracy and performance of the algorithm.

3.4.2 Genetic algorithm and pattern search

The genetic algorithm [103, 143, 52] belongs to a larger class of optimization search algorithms, known as evolutionary algorithms that work on the principle of natural selection. An initial population of the parameter θ is (randomly) chosen and the objective function is evaluated for each individual in the parameter population. The value of the objective function usually represents the fitness of the individual. A fixed number of fittest members of the parameter population are stochastically chosen to be evolved at the next iteration. The parameter is usually characterized by specific features or properties, which are then mutated or altered to obtain a new population or generation of the parameter at the next iteration. The process is repeated until a desired fitness or maximum value of the objective function is achieved or predefined maximum number of generations is reached.

The genetic algorithm is suitable for both constrained and unconstrained optimization problems. It is specifically used for discontinuous, non-differentiable stochastic functions to obtain a global extremum; however, it has a slow convergence rate and requires sufficiently large number of iterations to converge.

On the other hand, pattern search [20, 67, 146, 89] converges quickly to the solution. The search is initiated by evaluating the objective function at an initial value of the parameter. Then, successively neighboring parameter points are found for which the value of the objective function increases. The new parameter points are found using a mesh around the previous ones, and the search continues until a maximum number of iterations is reached or mesh size is too small.

3.4. OPTIMIZATION TECHNIQUES

Pattern search belongs to the family of direct search algorithms and is applied for discontinuous or non-differentiable objective functions.

We tested both the genetic algorithm and pattern search to optimize the log-likelihood function of our models. In the case of simple models having less than 4 parameters, the results were consistent with those obtained using the grid search method. However, the algorithms did not work for high-dimensional models. The involved computations makes it difficult for the algorithms to be run in parallel and in general, they require a huge amount of computational time and resources. Though these methods are certainly better and more precise than grid search algorithm, we were unable to obtain results using them for our models. Hence, our analysis is only based on the results obtained using grid search optimization algorithm.

Thus far, we discussed the procedure to obtain the predictions of a model given the data. We summarized the estimation method for model parameters and different optimization techniques that can be used for numerical estimation. We now focus on the visualization of our experimental data. We described the format of our collected data on the experiment in Section 3.1, but the question is how we could plot the data for analysis purposes. In the following section, we discuss the different types of *psychometric curves* we have used to analyze our data. Psychometric curves are frequently used in psychophysics to study the responses of subjects on an experiment. The predictions of a model are generated for these curves using model fitting techniques (discussed in Section 3.2). The error between a subject's actual psychometric curve and the model predicted curve provides a measure of

goodness of a model fit.

3.5 Psychometric curves

Psychometric curves [157, 77] are extensively used in psychophysics to represent the summary of subjects' responses on an experiment. They describe the behavior of subjects over a range of stimulus values. The curve is determined by the number of alternative choices in the task, for example, a psychophysical task can have binary choice, two-alternative forced choice (2AFC), or n-alternative choices. We are mainly interested in the curves having binary choices as our target detection task has a yes/no paradigm.

For our analysis purposes, we examine the behavior of subjects on our experiment using three different types of psychometric curves. We note that our experiment was divided into four sessions (details in Section 2.5) and each session was characterized by the unique value of correlation coefficient, ρ_s used to generate the stimuli. The purpose of our study is to determine whether subjects learn the structure present in the scenes and infer the correct correlation strength to make their decisions. Therefore, correlation coefficient, ρ_s serves as the physical stimulus parameter of our interest and we specify our results in terms of this parameter. We consider three types of psychometric curves based on different characterizations of stimuli as a function of ρ_s .

3.5.1 Types of psychometric curves

We consider the following three types of psychometric curves to analyze the responses of subjects on our target detection experiment described in Chapter 2:

- (I) **Hit and false-alarm rates:** In a target detection task, the probability of reporting "target present" when the target is present is known as the *hit rate* or *detection rate*, or *true positive rate*, and is denoted by $P(\hat{T} = 1|T = 1)$. Whereas, the probability of reporting "target present" when the target is absent is known as the *false-alarm rate* or *false-positive rate* and is denoted by $P(\hat{T} = 1|T = 0)$. The origin of these terminologies lies in the *signal detection theory* [56, 97, 158, 73, 101]. Further, we can obtain *miss rate* or *false negative rate* and *correct rejection rate* or *true negative rate* by subtracting the hit and false-alarm rates, respectively from 1.

We plot the hit and false-alarm rates as a function of correlation strength ρ_s in four different experimental conditions. We also compare the performance of subjects as the strength of ρ_s varies in the experiment.

- (II) We also plot the **proportion of "target present" responses as a function of minimum difference between the target and a distractor's orientation**. We plot these curves separately for the target-present and target-absent trials.

In a target-absent trial, as the distractor orientation gets closer to the target orientation, the difference between the two reduces, and it would be difficult for an observer to discriminate the distractor from the target. The distractor would appear as a target in such a case and the observer is more likely to

report "target present" even though there is no target. On the other hand, as the minimum difference increases, the distractors will have significantly different orientations than the target and the observer may be able to detect the absence of target easily.

While on a target-present trial, a comparatively small orientation difference between the target and a distractor could provide more evidence to the observer in responding "target present". This is because more than one object in the display would have orientation closer to the target and the observer could possibly make a decision based on either one of them. In the case of large minimum difference between the target and a distractor, the observer must make a decision based on the measurement of the target.

(III) Further, we analyze the behavior of subjects using a plot of **proportion of "target present" responses as a function of sample standard deviation between distractor orientations**. Again, we separate the data on target-present and target-absent trials. In terms of interpretation, these curves are closely related to type (II) curves with x-axis as the minimum orientation difference between target and a distractor. Also, the curve follows a similar shape as its counterpart in type (II) except in the case of perfect correlations. In the case of $\rho_s = 1$, all distractors are identical and hence the standard deviation between any pair of distractor orientation is equal. Thus, all data points lie in a single bin on target absent trials and we only obtain a single point in the plot. For instance, left panel in Figure 4.1(C) and other similar figures in Chapter 4 have such curves.

3.5.2 Predictions using synthetic data

We want to fit the psychometric curves of subjects on the experiment with a hypothesized model M . We use the maximum-likelihood estimate $\hat{\theta}$ of model parameter θ to generate the predictions of the model for a subject's data. We assume that $\hat{\theta}$ represents the subject's parameter and can be used to completely describe the behavior of the subject. We use this estimate to generate the stimuli and the corresponding hypothesized measurements of the subject. These measurements are then used to make hypothesized responses of the subject based on the decision variable $d_M(\mathbf{x})$ of the underlying model. This constitutes the model predicted data set for the subject and is also referred as the *synthetic* or *fake data set*. The synthetic data are generated to replicate the subject's behavior based on the assumptions of the model M . These data sets are then used to make model generated psychometric curves. The psychometric curves predicted by the model with the parameter estimate of $\hat{\theta}$ are compared with the actual psychometric curves of the subject on the experiment. If the model predicted curves are well fitted to the subject's true responses, then the model provides a good explanation of the behavior of the subject on the experiment. If the two curves are significantly different, then the underlying model lacks the assumptions that could reproduce the subject's responses.

It is possible that the responses of the subjects are best described by multiple models or behavior of different individuals is explained by different models. We find such issues in Chapters 4 and 5. But, we use summary statistics to draw

conclusions. We generate the model predictions for each subject and take an average over all subjects. These average model predictions are then matched with averaged true psychometric curves of the subjects. The results based on summary statistics may not reflect a complete picture in the event subjects follow different strategies or models to make decisions. However, it is difficult to build a mathematical model to test such a possibility.

3.5.3 Error measures

The model provides a prediction for the subject's data based on the maximum-likelihood estimate of parameter θ . The subject's psychometric curves are fitted using the model predicted curves. To quantify the difference between both curves, we measure two types of statistical errors: (i) the *root-mean-square error (RMSE)* and (ii) the *R^2 statistic or coefficient of determination*. RMSE is defined as the square root of the mean square error between the subject's data, y_i and the model predicted curves, \hat{y}_i :

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{D_N} (y_i - \hat{y}_i)^2}{D_N}}.$$

Here D_N denotes the total number of data points in the psychometric curve.

Statistic R^2 is another measure of determining how well the model fits to the

experimental data and it is defined as:

$$R^2 = 1 - \frac{\sum_{i=1}^{D_N} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{D_N} (y_i - \bar{y})^2},$$

where $\bar{y} = \frac{1}{D_N} \sum_{i=1}^{D_N} y_i$ is the sample mean of subject's data. The range of R^2 depends on the type of regression used; however, a negative value can occur in case a non-linear function is fitted to the data [32].

3.6 Model comparison

The fitting of model predicted psychometric curves to a subject's responses determines the goodness of the model to the experimental data and we measure this goodness in terms of the statistical error between the two curves. Frequently, we consider multiple models with different assumptions about the behavior of the subject. Models generally differ in terms of the assumptions about their parameters and the dimensionality of the parameter space. We fit each hypothesized model to the data using a similar fitting procedure. It is possible that more than one model provide a good explanation for the subject's behavior. They could have equally well predicted curves that match the experimental data and the error is comparable for both models.

How do we compare models to find the one that best describes the data? Psychometric curves do not help in discriminating two models if both predict similar

fitting curves for the data. Instead, we compare models on standard criteria that also weigh models based on their dimensionality. A general model with large number of parameters will always provide a fit that is at least as good as a special model obtained from the general case. This could mainly be because of additional free parameters in the general model. Therefore, we consider criteria that take this issue into account and penalize a model for the number of free parameters it has. In the following section, we briefly discuss three criteria that are commonly used to compare models in psychophysics. We also discuss the limitations we encountered while comparing our models (described in Chapter 5) based on these criteria.

3.6.1 Bayesian model comparison (BMC)

Bayesian model comparison (BMC) [96, 93, 155] is a fundamental method of model selection. The model with the highest posterior probability (probability of the model given the data) is selected. The posterior probability of a model M given the data can be computed using Bayes' theorem

$$p(M|\text{data}) = \frac{p(\text{data}|M)p(M)}{p(\text{data})}.$$

The ratio of posterior probabilities is computed for the two models M_1 and M_2 that need to be compared

$$\frac{p(M_1|\text{data})}{p(M_2|\text{data})} = \frac{p(\text{data}|M_1)p(M_1)}{p(\text{data}|M_2)p(M_2)}.$$

3.6. MODEL COMPARISON

An equal prior is chosen for both models, $p(M_1) = p(M_2) = \frac{1}{2}$ since no model is favored over the other. This reduces the comparison to the ratio of *model likelihoods*, $p(\text{data}|M)$. The ratio is also called a *Bayes' factor*. The *model likelihood* is computed [93, 98] by averaging the model likelihood under a hypothesized value of the parameter in the parameter space

$$p(\text{data}|M) = \int p(\text{data}|M, \theta)p(\theta|M)d\theta. \quad (3.4)$$

The above expression is rewritten in terms by taking logarithm

$$\log p(\text{data}|M) = \log L_M^* + \log \int e^{\log L_M(\theta) - \log L_M^*} p(\theta|M)d\theta. \quad (3.5)$$

Bayesian model comparison computed using the above equation is based on the entire model likelihood function $L_M(\theta)$, instead of only of its maximum value L_M^* . Further, it penalizes models for additional free parameters. Therefore, BMC is a principal method to compare models.

However, there are several issues in the practical implementation of BMC. Eq. (3.5) is based on the integral evaluation of the parameter likelihood function $L_M(\theta)$ over the entire parameter space. This integral does not have a closed form expression in case of analytically intractable model likelihood. Thus in practice, the integral is approximated by Riemann sums. The accuracy of such an approximation depends on many factors. For instance, the parameter grid spacing and number of samples used in the evaluation can potentially change the results to a large extent. Further, the sum may not converge to the true value and could suffer numerical issues. High-dimensional models are most sensitive to these problems

and precaution must be taken when using Bayesian model comparison for complex models.

Most of the models that we considered to explain our experimental data are high-dimensional in parameter space. Thus, we faced computational issues using BMC to compare our models. In general, the integral in Eq. (3.5) can be evaluated using *Markov Chain Monte Carlo (MCMC)* method [50, 121, 96, 107, 19]. However, we also encountered problems in implementing this method to apply BMC in our case. Lack of an analytical approximation for the model likelihood function, $L_M(\theta)$ created most of the numerical problems for us. Due to computational inefficiency of our models, we use criteria that use maximum model likelihoods to compare models. We describe two such comparison measures below.

3.6.2 Bayesian information criterion (BIC)

Bayesian information criterion (BIC) [131] is a model selection criterion that is based on the maximum value of model likelihoods L_M^* . It computes the goodness of a model from maximum model likelihood while penalizing the model for extra free parameters. The penalty term increases with the number of free parameters in the model. BIC is mainly computed using the following formula

$$\text{BIC} = -2L_M^* + f \log K, \quad (3.6)$$

where f is the total number of free parameters in the model and K represents the total number of experiment trials. The above formula is based on a Laplace's approximation of the integral in Eq. (3.5). Bhat et al. [18] have provided a detailed

derivation of the same.

The penalty term in BIC increases with increasing number of free parameters, f and it scales with the size of the data. Thus, it more strongly penalizes models with large number of free parameters as compared to the *Akaike information criterion (AIC)*. Therefore, the model with the lower value of BIC is preferred when using Bayesian information criterion to compare models. A lower value of BIC for a model implies fewer free parameters, better fit or both.

3.6.3 Akaike information criterion (AIC)

Similar to BMC and BIC, Akaike information criterion(AIC) [3, 21, 22, 27] is also a relative measure of a statistical model quality. It is closely related to BIC in terms of the criterion and the penalty term. However, the penalty term is weaker as compared to BIC. The following equation is generally used to compute Akaike information criterion:

$$\text{AIC} = -2L_M^* + 2f. \quad (3.7)$$

Under this criterion of model selection, the model with the minimum value of AIC is preferred over others and it also penalizes models for over fitting.

The derivation of the criterion in Eq. (3.7) can be found in [3]. It is mainly based on the information theoretic concept and computing the *Kullback-Leibler (KL) divergence* [85, 27, 19] between a model prediction and the data. The model with the minimum value of KL distance or having minimized information loss is selected.

We note that both BIC and AIC are only based on maximum model likelihoods.

Thus, this requires finding the maximum-likelihood estimates; however, we do not need to compute any integrals here. Therefore, to overcome our numerical difficulties, we compare our models using BIC and AIC in Chapter 5 and make conclusions based on the results obtained from these comparisons.

3.7 Discussion

In this chapter, we have presented a brief description of the tools that are frequently used in fitting models to the psychophysical experimental data. Estimating maximum-likelihood parameters of a model is a crucial step in order to fit the model to a subject's responses. However, possible numerical problems must be diagnosed that are specific to the task and the data. A suitable and efficient optimization algorithm must be applied to estimate model parameters. Lack of a better approximation of the underlying model likelihood or an inefficient method of obtaining this approximation can limit the application of an optimization algorithm to the data.

The parameter estimates are then used to generate the model predictions for the psychometric curves of the subject's responses. These curves are analyzed to determine the goodness of the model fit to the data and how well the model explains the behavioral response of the subject on the experiment. Eventually, different models are compared on an appropriate selection criterion to find the model that best describes the data.

Chapter 4

Data analysis I: model fitting

Recent experimental studies show that humans perform near-optimally in detecting a target among identically and independently oriented distractors [94, 98, 99]. Both homogeneous and heterogeneous distractors are extreme structural conditions and do not represent the possible realistic intermediate structure of natural scenes. It is unknown that how humans make decisions in response to weakly structured input in visual search tasks. Therefore, we explore whether humans can take into account partial stimulus correlations in their inference process.

We conducted a target detection experiment to examine the decisions of human observers in response to structured input. We interpolated between the extreme conditions of heterogeneous and homogeneous distractors, and introduced intermediate correlation strengths among distractor orientations. The correlation coefficient, ρ_s defined in the generative model in Section 2.1 controls the amount

of pairwise correlations between distractors. A high value of ρ_s implies more similar distractors and hence more structured visual input.

We consider several plausible models and rigorous model selection tools to find a model that best explains the data. We interpret the behavior of subjects on the experiment based on the assumptions of the best fitting model and its parameters. We wish to understand whether subjects take into account the strength and structure of stimulus correlations in their inference process. The correlation coefficient, ρ_s is an experimental parameter and we varied it to control the amount of statistical structure in visual scenes. We chose 4 different values of ρ_s - 0, $\frac{1}{3}$, $\frac{2}{3}$, and 1 to test in the experiment. The condition $\rho_s = 0$ corresponds to heterogeneous distractors and $\rho_s = 1$ to perfectly correlated (or homogeneous) distractors. The intermediate value of $\rho_s \in \{\frac{1}{3}, \frac{2}{3}\}$ interpolates between the two extreme conditions and introduces a weak structure in the scenes. The experiment was divided in four different sessions (details in Section 2.5.3) and a unique value of ρ_s was used to generate stimuli in each session. Therefore, each session is characterized by a unique statistical visual structure and represents a different condition on the experiment. We therefore refer to these sessions as *experimental conditions*, which differ by the experimental value of ρ_s used in that session.

Our goal is to examine the behavior of subjects on different experimental conditions and determine whether subjects use correct or incorrect assumption about the experimental value of ρ_s used in different conditions. Therefore, we consider ρ_s in our models and fit it per subject using maximum-likelihood estimation. In

addition to estimating ρ_s as a free parameter in a model, we also need to estimate the precision parameters which determine the subject's measurements. It is evident from Eq. (2.21) that precision (inverse of measurement variance, σ_i^2) of a measurement can greatly affect the accuracy of a subject's decision. But, we are unaware of the precision values with which subjects made their measurements. Therefore, we consider different possible assumptions about encoding precision of subjects in our models.

We consider a variety of models in this chapter that differ in the assumption about ρ_s and encoding precision. We consider a range of assumptions about both parameters in the models. Each model is individually tested on the experimental data and model predictions are generated for subjects' psychometric curves. We use the methods and procedures described in Chapter 3 for finding parameter estimates and fitting a model to the data. We measure the RMSE error and goodness of fit (R^2) between model predicted curves and the data. We conclude by discussing the need for comparing models.

4.1 Psychometric curves of subjects data

We first examine the experimentally obtained psychometric curves that are based on subjects' responses. We consider three types of psychometric curves: (I) hit and false-alarm rates, (II) proportion of "target present" responses as a function

of minimum difference between the target and any distractor, and (III) proportion of "target present" responses as a function of sample standard deviation between stimuli orientations. We separately analyze type (II) and (III) curves in target-present and target-absent trials. A detailed description of these curves is presented in Section 3.5.1. Figure 4.1 shows the mean responses of subjects on these different types of psychometric curves.

4.1.1 Hit and false-alarm rates

The left panel in Figure 4.1(A) shows the hit (black) and false-alarm rates (red curve) for average subject. The mean proportion of subjects responding "target present" on target-present and target-absent trials are plotted in each experimental condition ($\rho_s \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$). The hit and false-alarm rates show a constant trend in the first three experimental conditions, while they have an expected increase (hit rate) and decrease (false-alarm rate) in the case of homogeneous distractors. Increased pairwise correlations between distractor orientations result in more structure among the stimuli and would make it easier to single out the target, if present. Therefore, the increase in hit rate (or decrease in false rate) in presence of homogeneous distractors is due to the increased correlations that facilitate the target detection.

A similar trend is also seen in the right panel of Figure 4.1(A). The average subject performance is plotted in each experimental condition. The mean performance in first three conditions is very close to 60%, while in the case of $\rho_s = 1$,




Figure 4.1: **Psychometric curves based on the experimental data.** Throughout the chapter, the error bars indicate unit standard error mean (s.e.m). **(A) Hit and false-alarm rates (left), and performance (right).** (Left) Hit (black) and false-alarm (red) rates as a function of correlation strength, ρ_s used in the experimental conditions. Hit (or false-alarm) rate shows a near-constant behavior in the first three experimental conditions and a large increase (or decrease) for $\rho_s = 1$. (Right) Mean subject performance in the four experimental conditions. **(B) Minimum target-distractor orientation difference.** Proportion of "target present" responses as a function of minimum difference between the target and any distractor, separately for target-present (left) and target-absent (right) trials in the four experimental conditions. Each curve corresponds to an experimental condition with $\rho_s \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$. Subjects' responses have a very similar trend for $\rho_s = 0, \frac{1}{3}$, and $\frac{2}{3}$, while their behavior is different when $\rho_s = 1$. **(C) Sample standard deviation between distractor orientations.** Proportion "target present" responses on different experimental conditions as a function of sample standard deviation between distractor orientations, in target-present (left) and target-absent (right) trials. The curves are similar to those in (B) except for $\rho_s = 1$ in target-absent trials (right). Since all distractors are identical in a trial when $\rho_s = 1$, there is only a single data point for the sample standard deviation.

it shows a large increase to 72.7%. This suggests that subjects might be using the enhanced visual structure to improve their decisions. We will investigate this in further details below.

4.1.2 Minimum target-distractor orientation difference

A more detailed view of the data can be seen in Figure 4.1(B). The mean proportion of "target present" responses in all four experimental conditions are plotted as a function of minimum difference between the target and any distractor in both target-present (left) and target-absent (right) trials. These plots show that the proportion of subjects responding "target present" decreases as the minimum target-distractor orientation difference increases, both for target-present and target-absent trials. Such a behavior is expected, since a large difference between the target, and any distractor reflects more dissimilarity of the distractors from the target and hence it would be easier to perform the task. The curves corresponding to $\rho_s = 0, \frac{1}{3},$ and $\frac{2}{3}$ overlap, while the responses of subjects have a different behavior in the case of $\rho_s = 1$. This behavior is consistent with the trend seen in hit and false-alarm rates in (A). Also, the decrease in the proportion of "target present" responses at larger minimum target-distractor differences, is higher in target-absent trials for $\rho_s = 1$ as compared to other experimental conditions. This is because the distractors are identical when $\rho_s = 1$ and as the difference between target and distractors increases, it would become easier to determine whether all stimuli are same or there is an odd-ball (target) stimulus.

4.1.3 Sample standard deviation of distractor orientations

Another view of the data on different experimental conditions is presented in terms of proportion "target present" responses as a function of sample standard deviation of distractor orientations in Figure 4.1(C). Again, these are plotted separately for target-present (left) and target-absent (right) trials. These curves contain similar information as minimum target-distractor orientation difference plots in (B). Thus, the behavior on target-present trials is very similar to the left panel figure in (B). The proportion of "target present" responses decreases as the sample standard deviation of distractors increases and hence the distractors get more dissimilar. On target-absent trials, the sample standard deviation reduces to a single value in the case of homogeneous distractors and thus the curve has only one data point.

Figure 4.1 presents different psychometric curves to visualize the experimental data and each curve provides a different insight about subjects' behavior. Different plots suggest that subjects' responses are similar in experimental conditions with $\rho_s < 1$, while they do behave differently in the case of homogeneous distractors. We next describe different plausible models that could explain these observations and provide a good fit for the experimentally obtained psychometric curves.

4.2 Models

In the optimal-observer model derived in Section 2.2, we assume that the observer is aware of the correct generative model (Figure 2.1) and the associated statistical variables such as N, σ_s^2, ρ_s , and prior over T . However, we need to test whether human observers learn and use the correct generative model in their decisions. If they do not, their inferences are *suboptimal*. For instance, an observer may not use equal odds prior or the correct value of ρ_s . Therefore, we analyze how subjects made their decisions and what parameter values they used to make their responses. Specifically, we wish to determine what values of ρ_s subjects used to make their responses and how certain they were in making their measurements. We note that we can only determine this by fitting models to the data and making conclusions based on the best fitting model. However, it is always possible that there are better models and better explanations.

The subjects were pre-informed about the number of stimuli being $N = 4$ in the experiment. Further, we assume that subjects were able to infer the correct value of σ_s as 15° in the experiment and used it to make their decisions. Though, this may not necessarily be true. It is possible that subjects did not infer the value of σ_s correctly in the experiment and might have used an incorrect assumption about it. In that case, we would need models that incorporate the plausible assumptions about σ_s that observers could have used. Such models will have σ_s as a free parameter and we would need to test possible assumptions on it. For instance, subjects could use different values of σ_s across experimental sessions or

possibly they could use a distribution over the values of σ_s . We can easily see that this would add another dimensions of complexity to our existing complex models and would further make our computations intractable. Thus, we save some order of complexity in our models by assuming that subjects were able to correctly infer the true value of σ_s in the experiment.

We are mainly interested in determining whether subjects use correct or incorrect assumption about ρ_s in making decisions on different experimental conditions. We denote the true experimental value of ρ_s by $\rho_{s_{\text{true}}}$ and a subject's assumed value of ρ_s by $\rho_{s_{\text{assumed}}}$. Therefore, we want to determine whether $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$ or $\rho_{s_{\text{assumed}}} \neq \rho_{s_{\text{true}}}$ for a subject in each experimental condition. If the subject uses $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$ in all experimental conditions, the subject is *optimal* on the task. Otherwise, we refer to the condition $\rho_{s_{\text{assumed}}} \neq \rho_{s_{\text{true}}}$ as *suboptimal condition*.

To answer this question, we consider several models. We categorize these models based on the assumptions about ρ_s and the encoding precision:

1. **Assumption about ρ_s :** we explored whether subjects use the correct ($\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) or incorrect ($\rho_{s_{\text{assumed}}} \neq \rho_{s_{\text{true}}}$) assumption about ρ_s in the generative model.
2. **Encoding precision:** we do not know how subjects made their measurements. In particular, what precision values they used for encoding stimuli. We need to determine whether subjects encode all stimuli with equal precision or they use varying precision for stimuli across trials. Therefore, we test models which assume that subjects either encode stimuli with equal or

variable precision.

In the following sections, we discuss the above two categories of models in details and also their further division into sub-categories.

4.2.1 Assumptions about ρ_s

An ideal observer uses the correct value of ρ_s when inferring target presence. However, it is possible that an observer may not be aware of the correct correlation strength. To test which case is more likely, we consider models with ρ_s as a free parameter. Specifically, we tested three main assumptions about ρ_s :

- (a) $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$: we assume that the observer uses the correct value of correlation strength in all experimental conditions to infer target presence, that is, $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}} \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$.
- (b) $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{constant}}}$: in this condition, we consider that the observer assumes that the correlation strength among distractor orientations is constant across all conditions. In particular, we check whether
 - (i) the observer completely ignores the information about structural correlations in all experimental conditions and uses $\rho_{s_{\text{constant}}} = 0$. This would mean that observer uses $(0, 0, 0, 0)$ as correlation values in making decision on all experimental conditions.
 - (ii) the observer may also use any other value between 0 and 1 as $\rho_{s_{\text{constant}}}$ in making decision. We thus let $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{constant}}}$ to be a free parameter in

the model.

(c) $\rho_{s_{\text{assumed}}} \neq \rho_{s_{\text{true}}}$: we further allow the possibility that the observer may use different, possibly incorrect correlation strengths across experimental conditions. We let $\rho_{s_{\text{assumed}}}$ to be a free parameter in the model and check for the following possibilities:

- (i) the observer considers the first three conditions identically and thus uses an equal correlation strength in making decision on these conditions, while assume a different correlation strength in the fourth experimental condition. In such a case, we fit $\rho_{s_{\text{assumed}}}$ as a constant free parameter in the first three conditions and as another free parameter in the fourth experimental condition, *i.e.*, $(\alpha, \alpha, \alpha, \beta)$.
- (ii) the observer uses an incorrect assumption about ρ_s in all experimental conditions. We thus fit $\rho_{s_{\text{assumed}}}$ per condition in this model, *i.e.*, $\rho_{s_{\text{assumed}}}$ has following form $(\alpha, \beta, \gamma, \delta)$. We note that this is the most general assumption about ρ_s and all above models are special cases of this model.

Here α, β, γ , and δ represent free parameters of the model and these are fitted per model for each subject. We also note that $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{constant}}}$ and $\rho_{s_{\text{assumed}}} \neq \rho_{s_{\text{true}}}$ are suboptimal behavioral conditions.

4.2.2 Encoding precision assumptions

Along with the assumption about ρ_s , we also consider different assumptions about encoding precision in our models. Since the precision of measurements greatly impacts the accuracy of decisions, we wish to determine which assumption is most consistent with the responses of the subjects. Signal detection models have typically assumed that encoding precision is constant across stimuli and trials at a given set size [110, 135, 164, 159]. However, recent experimental studies show that observers' measurements are of variable precision. That is, the encoding precision varies across stimuli and trials [150, 46, 98, 99, 75, 76, 137, 139]. This variability could be attributed to the attentional fluctuations or other factors.

We thus consider both possible models of the precision of measurements. In the first type of models, we assume that the observer assigns equal precision to all stimulus measurements. We denote the precision of the i^{th} measurement by $J_i = \frac{1}{\sigma_i^2}$. Under the equal precision assumption, we assume that J_i is constant across stimuli, that is, $J_i = J$ for all $i = 1, 2, \dots, N$. The constant J is a free parameter in the model. These models are known as **equal precision (EP) models**. We assume that the precision is constant across trials in the same experimental condition; however, it may vary across different experimental conditions. We thus test both possibilities in our models by assuming J as constant and variable across experimental conditions. We, therefore, consider two types of EP models with J as a single constant parameter in all conditions and as a varying parameter across the four experimental conditions.

In the second type of precision models, we assume encoding precision varies randomly across stimuli and trials. These models are known as **variable precision (VP) models**. We denote the precision variable by a vector $\mathbf{J} = (J_1, J_2, \dots, J_N)$, where J_i corresponds to the precision of the i^{th} measurement. In the VP models, we assume that the precision variable \mathbf{J} with which stimuli are encoded is a random variable. To model such variability, we assume \mathbf{J} follows a gamma distribution with mean \bar{J} and a scale parameter τ . We sample the precision randomly for each stimulus on each trial. Therefore, the measurement is described by a doubly stochastic process, $(\bar{J}, \tau) \rightarrow \mathbf{J} \rightarrow \mathbf{x}$ [150]. Thus, the precision determines the distribution of stimulus estimate, but is itself also a random variable.

Why do we choose the gamma distribution to model the variability in precision? The proper choice of a distribution for modeling variable precision would require the marginalization over all possible ways to implement this variability. Though a full marginalization seems impossible, but the success of the VP concept can be assessed how well it performs under various specific alternative distributions. Van den Berg et al. [149] have implemented and tested VP models with many other alternatives such as log-normal, Weibull, and log-uniform distributions. They have found that the results are consistent under changes in the assumed distribution over precision. Therefore, we consider that our results would be robust under the choice of a distribution. The gamma distribution is a two-parameter family of continuous, unimodal distributions on the positive real line and has been successfully used for modeling variable precision. We thus consider the same choice for our models.

In our VP models, both \bar{J} and τ are free parameters. We assumed the scale τ to be constant across experimental conditions. Similar to EP models, we allow \bar{J} to be constant or varying across experimental conditions in the models.

4.2.3 Summary of models

We test the following two assumptions about encoding precision in EP (with precision parameter J) and VP (mean precision parameter \bar{J}) models:

- precision J (or \bar{J} in VP) is independent of the experimental conditions (ρ_s -independent),
- precision J (or \bar{J} in VP) vary across experimental conditions and is thus ρ_s -dependent.

In addition, we have following model variants based on the assumption about ρ_s for each category of EP and VP models:

- (i) $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$, the optimal model,
- (ii) $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$, *i.e.*, no correlations model,
- (iii) $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$, constant correlations model,
- (iv) $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$, constant in first three conditions and different in $\rho_s = 1$ condition, and
- (v) $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$, the most flexible model in terms of ρ_s .

We test each of the above assumption about ρ_s for each possible combination of EP, and VP models. Therefore, in total, we consider $2 \times 2 \times 5 = 20$ models. In addition to the assumptions about ρ_s and encoding precision, we also consider prior for T as a free parameter in our models. We assume this parameter to be constant across experimental conditions.

Table 4.1 gives a detailed summary of the models and their number of parameters. For reference convenience, we number models in each precision category, from EP1 to EP10 and VP1 to VP10. We note that most of the models are high-dimensional (with parameters > 3) and these parameters are associated with different correlation (experimental) conditions.

Precision	Assumption about J or \bar{J}	Model No.	Assumption about ρ_s	No. of free parameters
EP	J condition-independent	EP1	$\rho_{s_{\text{assumed}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$	2
		EP2	$\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$	2
		EP3	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$	3
		EP4	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$	4
		EP5	$\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$	6
EP	J condition-dependent	EP6	$\rho_{s_{\text{assumed}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$	5
		EP7	$\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$	5
		EP8	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$	6
		EP9	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$	7
		EP10	$\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$	9
VP	\bar{J} condition-independent	VP1	$\rho_{s_{\text{assumed}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$	3
		VP2	$\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$	3
		VP3	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$	4
		VP4	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$	5
		VP5	$\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$	7
VP	\bar{J} condition-dependent	VP6	$\rho_{s_{\text{assumed}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$	6
		VP7	$\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$	6
		VP8	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$	7
		VP9	$\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$	8
		VP10	$\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$	10

Table 4.1: **List of models fitted to the experimental data on the target detection task.** Description of models fitted to the experimental data with different assumptions about encoding precision parameter and correlation coefficient. The number of free parameters per model is also listed.

4.3. EQUAL PRECISION MODELS

Besides these models, there could be many other possible assumptions about ρ_s such as $\rho_{s_{\text{assumed}}} = (1, 1, 1, 1)$, $\rho_{s_{\text{assumed}}} = (0, 0, 0, 1)$, $\rho_{s_{\text{assumed}}} = (0, 0, 0, \alpha)$, and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, 1)$. We also test these possibilities in our models; however, the model fits in these cases are either worse or comparable to other general models we have considered above. Thus, we do not include these models in our analysis below.

We now present the model fits to the experimental data for each of the model listed in Table 4.1. Each model fit is obtained by generating synthetic data based on the maximum-likelihood estimates of the model parameters for each subject using the same number of trials as in the subject's data (averaged over 100 runs). The model predictions for different psychometric curves are obtained based on these synthetic data.

4.3 Equal precision models

We first examine the fits of the EP models to the subjects' data. We present the predictions of all EP models for the psychometric curves and discuss the consequences.

4.3.1 Condition-independent precision J

We consider precision to be constant across experimental conditions and test different assumptions about ρ_s . We check the model fits of EP1 to EP5 here. These

4.3. EQUAL PRECISION MODELS

models have a common assumption about precision, but they differ in the assumption about $\rho_{s_{\text{assumed}}}$ and hence have varying number of free parameters. In addition, they have a common assumption about the prior of T . In the figures that follow, the shaded areas represent the fits of the model to the data. The model predictions for each subject are individually obtained and averaged over subjects.

Figure 4.2 shows the fits of EP1 model for the data. This model assumes that subjects use correct assumption about the correlation strength, ρ_s in their inference process. The fits for the hit rate, false-alarm rate, and performance illustrate a good match (RMS errors of 0.036 and 0.011, and R^2 values are 0.94 and 0.96, respectively) between the model predictions and the data. Also, the model predictions have a close agreement to the data for the psychometric curves in Figure 4.2(B) with small errors (RMSE errors are 0.05, 0.05, 0.06, and 0.06 in the four experimental conditions, while R^2 values equal to 0.94, 0.92, 0.89, and 0.91, respectively). However, when we analyze the model fits based on sample standard deviation plots, the model predictions fail badly and show large deviations (Figure 4.2(C)). Specifically, the predictions are worse in the cases of $\rho_s = 0$ and $\rho_s = \frac{1}{3}$ (RMSE equal to 0.095 and 0.075; R^2 equal to 0.3 and 0.46, respectively). These curves clearly suggest that subjects do not use this model assumptions in their inference process.

We see a far more worse trend (RMS errors of more than 0.6 and R^2 values are even negative for some curves) in the fits of EP2 model in Figure 4.3. None of the subjects' curves are predicted by the assumption of zero correlations in this model. The model predictions are worse in the case of homogeneous distractors,




Figure 4.2: **EP model 1** (J condition-independent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. Throughout the chapter, the shaded areas show the fit of the model and the circles are averaged subject responses. The error bars and shaded areas represent unit standard error of the mean for subjects' data and model fits, respectively. RMS errors and goodness of fit, R^2 values between the data and model predictions are indicated by blue and green numbers in the plots, respectively. **(A) Hit and false-alarm rates.** (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . **(B) Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). **(C) Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).




Figure 4.3: EP model 2 (J condition-independent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.3. EQUAL PRECISION MODELS

$\rho_s = 1$ (Figure 4.3(C)) with the RMS errors of 0.2 and R^2 having a large negative value. The worse fits of this model in all conditions clearly indicate that subjects might be using correlations in making their decisions. However, we still need to investigate what correlation values they use in such a case.

Next, we compare the fits of EP3 model in Figure 4.4 that has the assumption of constant $\rho_{s_{\text{assumed}}}$ across experimental conditions. The model has better predictions for the data in the first three experimental conditions (RMSE < 0.07 and $R^2 > 0.64$, respectively) as compared to the condition of $\rho_s = 1$. Though the fits are not perfect even when $\rho_s = \frac{1}{3}$ and $\rho_s = \frac{2}{3}$. Particularly, the predictions in Figure 4.4(C) show disagreement between model fits and the data points at large values of sample standard deviation between distractor orientations. The model completely fails to explain the behavior of subjects on the homogeneous distractors condition having RMS errors of 0.1 and R^2 of 0.64 in the minimum target-distractor orientation plots (Figure 4.4(B)), and RMS errors of 0.15 and R^2 of -0.21 in the case of sample standard deviation plots (Figure 4.4(C)). These model fittings with the constant assumption of ρ_s suggest that perhaps subjects treat the case of $\rho_s = 1$ differently from the other experimental conditions and they might be using different inference processes in the case of perfectly structured inputs. This seems to be consistent with the responses of subjects observed in Figure 4.1. However, since this model could not completely explain the behavior even in the cases of $\rho_s < 1$, it is difficult to support this hypothesis using the predictions of the EP3 model.

4.3. EQUAL PRECISION MODELS




Figure 4.4: EP model 3 (J condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).




Figure 4.5: EP model 4 (J condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.3. EQUAL PRECISION MODELS

To test our hypothesis about whether subjects treat the homogeneous condition differently from other experimental conditions, we check the predictions of the EP4 model in Figure 4.5. EP4 model allows the possibility of having $\rho_{s_{\text{assumed}}}$ as a free parameter in the first three conditions and a different value in the fourth condition along with the assumption of constant precision parameter J . That is, $\rho_{s_{\text{assumed}}}$ has the form of $(\alpha, \alpha, \alpha, \beta)$, which implies that the observer uses a constant value of $\rho_{s_{\text{assumed}}}$ in the first three experimental conditions, but treats the fourth condition differently (when $\beta \neq \alpha$ for the observer). If, for an observer, β is equal to α , the fits of the EP4 model reduce to those of EP3 (with $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) for that observer's data.

The EP4 model has better predictions for the subjects' behavior on conditions with $\rho_s < 1$. Specifically, the hit and false-alarm rates (left), and performance curve (right) are closely fitted in Figure 4.5(A), with RMS errors of 0.031 and 0.023 (R^2 of 0.95 and 0.83), respectively. Further, the shaded curves of model are well aligned (RMS errors of 0.04 and $R^2 \geq 0.94$) with the data points in the minimum target-distractor orientation plots in Figure 4.5(B) except in the case of $\rho_s = 1$, where it shows some deviations in the target-present trials (fourth column). But, the comparison based on the sample standard deviation figures (Figure 4.5(C)) indicate that this model also fails to capture the responses of subjects on the perfect correlation condition. The model does have better fits for the data as compared to EP3 model in Figure 4.4(C, fourth column); however, they still have large RMS errors of 0.076 and low R^2 value of 0.63. In addition, the model predictions fail at some data points in the first three experimental conditions in Figure 4.5(C).




Figure 4.6: EP model 5 (J condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.3. EQUAL PRECISION MODELS

Therefore, the models EP1 to EP4 still lack the assumptions to explain the behavior of subjects on the target detection task. In order to further explore how subjects made their decisions, we test the most flexible assumption about $\rho_{s_{\text{assumed}}}$ with constant precision J across experimental conditions in EP models. We let $\rho_{s_{\text{assumed}}}$ be a free parameter and fit it per condition having the form of $(\alpha, \beta, \gamma, \delta)$. This assumption is included in our EP5 model and it represents a general model for the assumption on $\rho_{s_{\text{assumed}}}$. The model fitting curves are shown in Figure 4.6. Analyzing the model fits for each psychometric curve, we find that the predictions of this model are very close to those of EP4 model in Figure 4.5. Further, the magnitudes of RMS errors and the goodness of fit, R^2 are similar in both cases. This suggests that both models are close in their predictions for the data; however, none of them provide a good explanation for the data in the perfect correlation experimental condition.

Since none of the models EP1 to EP5 could provide a good fit for the experimental data, it is difficult to conclude anything about the behavior of subjects based on these models. However, these models were based on different assumptions about $\rho_{s_{\text{assumed}}}$, including the most general one but having the common hypothesis about constant J across experimental conditions. The mis-fit of all these models to the data suggests that in addition to $\rho_{s_{\text{assumed}}}$, the assumption of precision J is also important and that plays a crucial role in determining the predictions of the model for the data. Thus, we consider other models to test the assumption whether subjects have varying precision J across experimental conditions. To be consistent, we test similar assumptions about $\rho_{s_{\text{assumed}}}$ as in models EP1 to EP5 in

the following section.

4.3.2 Condition-dependent precision J

We now consider models EP6 to EP10, which have a common assumption about the precision J constant in a particular experimental condition, but varying across conditions. We begin with testing the hypothesis of $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$ in EP model 6.

We find in Figure 4.7(A) and (B) that the model has relatively good fit ($R^2 \geq 0.89$) for the data, as compared to models EP1 to EP5, that assume J to be constant across experimental conditions. But, the model predictions are worse for all conditions in the sample standard deviation plots (Figure 4.7(C)). The RMS errors range from a minimum of 0.059 ($\rho_s = \frac{1}{3}$) to 0.082 ($\rho_s = 0$), while R^2 values are of the orders of 0.51 (when $\rho_s = 0$) to 0.57 ($\rho_s = 1$). Therefore, sample standard deviation plots provide us additional information about the model fits in addition to the minimum target-distractor orientation plots which were used for analysis in studies done by Mazyar et al. [98, 99]. Based on the poor performance of this model in Figure 4.7(C), we reject this model and test other assumptions about the $\rho_{s_{\text{assumed}}}$ for subjects' responses.

We next consider the fits of EP7 model in Figure 4.8. This model assumes that observers do not learn any structural correlations present in visual scenes and consider all conditions as heterogeneous displays ($\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$). Considering only the curves in Figure 4.8(A), the model seems to be predicting well for




Figure 4.7: EP model 6 (J condition-dependent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).




Figure 4.8: EP model 7 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

the data with small RMS errors of 0.035 (left panel) and 0.025 (right panel). However, when we consider the fits for the minimum target-distractor orientations in Figure 4.8(B) and the sample standard deviation plots in Figure 4.8(C), we see huge errors (RMSE of the order of 0.06 or more) between the model predictions and the data. This clearly indicates that this model do not explain the responses of the subjects in the experiment. Perhaps subjects do take correlations into account while making their decisions. But, we are unaware of the values of the structural correlations they use.

Thus, we test the assumption of constant correlations in the experiment in EP model 8. The model predictions and fits are shown in Figure 4.9. Analyzing the curves in (B) and (C), we find that the model does not reproduce the subjects' behavior on the experiment. It particularly fails in the condition of $\rho_s = 1$ with RMS error of more than 0.1 and relatively poor R^2 value. We also find that the fitting of psychometric curve in the experimental condition of $\rho_s = \frac{1}{3}$ in Figure 4.9 (B, second column) is worse as compared to EP7 model in Figure 4.8(B, second panel). This seems odd, since EP7 model is a special case of EP8 model with $\alpha = 0$ in all conditions. This might be explained using the computational inefficiency we faced in obtaining the fits of our models. We obtained the maximum-likelihood model parameters using the exhaustive grid search method (Section 3.4.1) and such an estimation depends on the convergence of the likelihood function. If the likelihood function does not converge, the estimates may not represent the global maximum-likelihood parameter estimates. But, overall the model has good




Figure 4.9: EP model 8 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

predictions as compared to EP7 model, particularly comparing the fits in Figure 4.9(C) to the ones in Figure 4.8(C).

We continue our analysis with EP9 model having condition-dependent precision parameter J . The model provides a better picture of the data as seen in Figure 4.10. We find better fits for the psychometric curves in (A) and (B). However, we continue to find model deviations for the sample standard deviation plots in Figure 4.10(C). In particular, the model fails to account for the data in the homogeneous condition. Though compared to other EP models (EP1 to EP8) discussed so far, the RMS error is found to be small in this case at 0.058 (comparable in case of EP1) and the model makes good predictions for the data but it still does not completely explain the subjects' behavior based on the sample standard deviation plots.

Finally, we test our last EP model listed in Table 4.1, which is EP10 model. The EP model has the most flexible assumption about $\rho_{s_{\text{assumed}}}$ and precision J . We compare the fits of this model in Figure 4.11. Like many other EP models, the model does extremely well in the first three conditions, but fails to explain the behavior on the homogeneous condition. Specifically, the poor model fit is continuously seen in the sample standard deviation plot for $\rho_s = 1$ (Figure 4.11(C) here). The hit and false-alarm rates (Figure 4.11(A)) along with the psychometric curves of minimum target-distractor orientations (Figure 4.11(B)) are well predicted by the model assumptions and have R^2 values of more than 0.9. Further, the model fits to the data are acceptable in the first three experimental conditions in Figure 4.11(C) with small RMS errors of up to 0.044. But, a worse model fit (RMSE




Figure 4.10: EP model 9 (J condition-dependent and $\rho_{s, \text{assumed}} = (\alpha, \alpha, \alpha, \beta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).




Figure 4.11: EP model 10 (J condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

equal to 0.064) in the condition of $\rho_s = 1$ allows us to reject this model and explore variable precision models.

The motivation to test VP models is to find a model that could possibly explain the subjects' behavior equally well across all experimental conditions and in terms of all psychometric curves. We thus now analyze the fits of VP models for the experimental data.

4.4 Variable precision models

Since equal precision models failed to explain the behavior of subjects' on the experiment, we further investigate if models with variable precision can provide an insight into how subjects inferred their responses. Variable precision models have been successfully used to explain human decisions on visual search task [98, 99] with homogeneous and heterogeneous distractors. We thus test if varying precision can be a key factor in interpreting the decisions of subjects in the target detection experiment here. Similar to EP models, we consider both possibilities of the mean precision \bar{j} of the gamma distribution to be constant or varying across experimental conditions. We discuss the fits for both categories of models below.

4.4.1 Condition-independent mean precision \bar{j}

We first examine the variable precision models with mean \bar{j} assumed to be constant across experimental conditions. We also assume the scale parameter τ of the

4.4. VARIABLE PRECISION MODELS

gamma distribution to be constant across experimental conditions and is another free parameter for the precision in the model. We check for the five possibilities for the $\rho_{s_{\text{assumed}}}$ in VP1 to VP5 models (described in Table 4.1) and analyze the resulting fits to find the best model if there is any.

Figure 4.12 shows the fitting of VP1 model to the data. We see the model explains the subjects' responses in the first three experimental conditions except for minor deviations in the target-present trials when $\rho_s = \frac{1}{3}$ in Figure 4.12(B). However, like many EP models, the model fails to account for the behavior in the perfect correlation condition. Specifically, the model predicts a lower proportion of "target-present" responses as compared to the subjects' data in the sample standard deviation plot in Figure 4.12(C, fourth panel). The RMS error is equal to 0.077, while the goodness of fit, R^2 is measured at 0.56 in this case. Therefore, with this variable precision model having correct assumption about the generative model, the data could not be explained well in the homogeneous condition.

We further check the predictions of VP2 model for the data in Figure 4.13. Very similar to VP1 model in Figure 4.12, the predictions of VP2 model completely fail to interpret the behavior in the experimental condition with $\rho_s = 1$. Moreover, the predictions are worse in this case as compared to VP1 model, as the RMS error increases to 0.15 (in Figure 4.13(C, fourth panel)) as compared to 0.077 in VP1 model (Figure 4.12(C, fourth panel)) and 0.1 (Figure 4.13(B, fourth panel)) from 0.08 (Figure 4.12(B, fourth panel)) in the minimum target-distractor orientation plots. However, except for the condition of $\rho_s = 1$, the data points are well matched with the model predictions. For instance, the goodness of fit




Figure 4.12: VP model 1 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.4. VARIABLE PRECISION MODELS




Figure 4.13: VP model 2 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

varies from 0.93 to 0.98 in the case of minimum target-distractor orientation plots in Figure 4.13(B, first, second, and third panels) and 0.9 to 0.97 in sample standard deviation plots in Figure 4.13(C, first, second, and third panels). This suggests that there is something really particular about the homogeneous condition that the above discussed models are missing and subjects do treat the condition with higher statistical structure differently as compared to the partial correlation conditions.

Further, we analyze the fits of models VP3, VP4, and VP5 in Figures 4.14, 4.15, and 4.16, respectively. We find that all these models have a similar trend in the first three conditions and completely fail to predict the behavior on the homogeneous condition ($\rho_s = 1$). Though these models differ in their assumptions about $\rho_{s_{\text{assumed}}}$, the predictions are quite similar in terms of model generated psychometric curves. Model VP5 has the most flexible assumption about $\rho_{s_{\text{assumed}}}$ in this category of models and we notice that this model also shows large RMS errors of 0.091 in Figure 4.16(C, fourth panel) and a poor fit of R^2 equal to 0.43.

Thus, none of the VP models in this category could provide suitable explanation for subjects' responses. All models VP1 to VP5 show good agreement with the data in the first three experimental conditions (with small RMS errors), while fail to incorporate the decision strategies of subjects when the distractors are perfectly correlated. The assumption of variable precision with constant \bar{J} across conditions did not improve the predictions of the models. This suggests that perhaps subjects' precision not only vary across stimuli and trials, but may also be varying across experimental conditions. Or, in the worst case, subjects might be using




Figure 4.14: VP model 3 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.4. VARIABLE PRECISION MODELS




Figure 4.15: VP model 4 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$ fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.4. VARIABLE PRECISION MODELS




Figure 4.16: VP model 5 (\bar{J} condition-independent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

some other assumptions about the generative model that our discussed models (EP1 to EP10, and VP1 to VP5) fail to incorporate.

Therefore, we finally test whether adding the assumption about varying \bar{J} across conditions in variable precision models could provide a better explanation for the subjects' behavior on the experiment, specifically in the case of homogeneous condition.

4.4.2 Condition-dependent mean precision \bar{J}

In this section, we analyze the predictions of variable precision models that assume mean precision parameter \bar{J} to be variable across experimental conditions and hence we fit \bar{J} per condition as a free parameter in the model for each subject. Models VP6 to VP10 described in Table 4.1 include this assumption about \bar{J} .

We examine the fits of VP6 model in Figure 4.17. We find that model predictions in this case are in close agreement to the data in all the conditions and for all types of psychometric curves. The goodness of fit measured in terms of R^2 is above 0.9 for all curves, except for the minimum target-distractor plot in Figure 4.17(B, fourth panel) where it is 0.88, as the model misses few data points. But overall, we see a well match between the model predictions and our data.

Thus, the question arises: do the subjects follow this model in their inference process? We note that VP6 model assumes that $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$ with varying \bar{J} across experimental conditions. The good fits of this model indicate that subjects were able to infer the correct information about the generative model and use it




Figure 4.17: VP model 6 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.4. VARIABLE PRECISION MODELS

to make their responses. But, we must investigate other models with alternative assumptions about $\rho_{s_{\text{assumed}}}$ before making a conclusion here. We need to examine how other models perform under the assumption of varying \bar{J} between experimental conditions.

To this end, we study the fitting of models VP7, VP8, VP9, and VP10 in Figures 4.18, 4.19, 4.20, and 4.21, respectively. We find that except for model VP7, all other models are equally good in their fits to the data. Model VP7 assumes that observers do not use any structural information in their decisions and Figure 4.18(C, fourth panel) shows poor fits of the predictions of this model for the data having RMS error of 0.11. Based on this model, we can assert that subjects do use stimulus correlations in making their responses and specifically the behavior of subjects, is remarkably different in the homogeneous condition as compared to the partial correlation conditions.

This hypothesis is well tested with model VP9 that assumes such a structural form of $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$. Figure 4.20 illustrates a good match between the model-generated curves and the psychometric curves based on subjects' data. The goodness of fit R^2 ranges from 0.91 to 1 for different curves and establishes that the model is successful in reproducing the responses of subjects on the experiment.

Further, the most general model VP10 perform equally well as models VP6, VP8, and VP9 in generating the predictions for the subjects' data. However, this model has more numbers of free parameters, as we allow the possibility of $\rho_{s_{\text{assumed}}}$ fitted per condition. The magnitude of RMS errors for this model is comparable with other models in this category (except model VP7). This indicates that this is




Figure 4.18: VP model 7 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).




Figure 4.19: VP model 8 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

4.4. VARIABLE PRECISION MODELS




Figure 4.20: VP model 9 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$) fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) **Minimum target-distractor orientation difference.** Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) **Sample standard deviation of distractor orientations.** Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).




Figure 4.21: VP model 10 (\bar{J} condition-dependent and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$ fits for the data. (A) Hit and false-alarm rates. (Left) Hit and false-alarm rates as a function of correlation strength, ρ_s used in the experimental conditions. (Right) Performance as a function of correlation strength, ρ_s . (B) Minimum target-distractor orientation difference. Proportion "target present" responses as a function of minimum target-distractor orientation difference, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns). (C) Sample standard deviation of distractor orientations. Proportion "target present" responses as a function of sample standard deviation of distractor orientations, separately for target-present (gray) and target-absent (red) trials in each experimental condition (columns).

the best possible fitting we can obtain for our data given these models.

However, we still have minor deviation between the model and the data at one particular data point, the last data point in target-absent trials (red curve) in the sample standard deviation plot for $\rho_s = \frac{1}{3}$ in Figure 4.21(C, third panel). This point is nearly missed by predictions of other VP models in this category (compare third panel in row (C) of Figures 4.17, 4.19, 4.20, and 4.21). It indicates that probably this is the best possible fitting we could obtain for the experimental data using these models. There may be something special about the characteristic of that point and we may be missing some particular assumptions in our models that result in the deviation of models at that point. We may need to put additional assumptions in the model to incorporate the fit at that particular point in Figure 4.21(C, third panel). However, we are unaware of the mathematical complexities it will introduce in our existing high-dimensional models. Further, model VP9 at least does a better job in predicting the behavior at this last data point on target-absent trials in Figure 4.20(C, third panel). Therefore, for our purposes, we limit ourselves to this level of fitting for the data and only focus on analyzing and comparing these models here.

4.5 Need for model comparison

In earlier sections, we analyzed the fits of equal and variable precision models listed in Table 4.1. We found that none of the EP models could account for the behavior of subjects on the target detection experiment (described in Chapter 2). The

4.5. NEED FOR MODEL COMPARISON

variable precision models with the assumption of mean precision \bar{J} constant across experimental conditions were also unsuccessful in reproducing the responses in the data. But VP models with variable \bar{J} (except model VP7) provided good predictions for the data in terms of low RMS errors and high R^2 values. However, we obtained more than one model that had equally well fits for the data. Specifically, models VP6, VP8, VP9, and VP10 all provided equally well predictions for the data. But, all these models have different underlying assumptions about $\rho_{s_{\text{assumed}}}$.

Therefore, we encounter the issue of discriminating multiple models that generate equally well predictions for the data. The question is how do we compare these models, since each model has different number of free parameters, and different assumptions about those parameters. Further, how do we find the model that provides the best explanation for the data. We thus use the model comparison techniques discussed in Section 3.6 to compare different models, and find the best fitting model. We present the model comparison results for our models in the following chapter.

Chapter 5

Data analysis II: model comparison

The purpose of the experimental study described in Chapter 2 was to determine whether human observers take stimulus correlations into account in a target detection task. We investigate whether observers are optimal in integrating the structural information in their decisions. If they are not, we wish to determine the alternative suboptimal processes observers use in their inference process.

To this end, we analyzed our experimental data in Chapter 4 using a variety of models. We considered different assumptions about encoding precision and correlation strength, $\rho_{s_{\text{assumed}}}$ in our models. We found that models with equal precision assumption did not fit well to the data; whereas the variable precision models provided better predictions for the data. We used gamma distribution to model the variability of precision across trials and stimuli in variable precision models.

We also tested the possibility if the mean precision parameter of the precision distribution can vary across experimental conditions. We observed that variable precision models with condition-dependent mean precision provided better predictions for the data as compared to models with condition-independence assumption. Moreover, there were multiple VP models (VP6, VP8, VP9, and VP10) that generated equally matched predictions for the subjects' data. Given the choice between different models (Table 4.1), we wish to select a model that explains the data most parsimoniously.

We continue our analysis from Chapter 4 here and compare models to find the best fitting model. We combine multiple models in a group or family that share a common characteristic and then find the best model in that group based on Bayesian or Akaike information criteria. We make our conclusions based on the hypothesis of the selected models. We use these comparisons to conclude why certain models are better than others. We also analyze the maximum-likelihood parameter estimates of the best fitting models and explain the behavior of subjects based on the estimated values of the model parameters.

5.1 Model families

We first combine our models into different families based on their assumptions about a certain parameter. We note that the models listed in Table 4.1 have varying number of parameters and they have different assumptions about those parameters. Some models share a common assumption about a particular parameter,

5.1. MODEL FAMILIES

while they differ about others. For example, models EP1 to EP10 assume equal encoding precision; however, they have varying assumptions about the correlation coefficient, $\rho_{s_{\text{assumed}}}$. Therefore, these models can be grouped into one family based on their common assumption about precision. Similarly, variable precision models (VP1 to VP10) represent another family of models based on the precision assumption. In this case, "precision" represents a common model characteristic or a *factor*, while equal and variable precision are two model families in this factor. All models in a family or a *level*, have a common assumption about the factor or characteristic it belongs to.

Similarly, we identify other factors and their respective model families in our list of models. We combine our models based on following three factors: (1) precision, (2) dependence of precision parameter, and (3) the assumption about $\rho_{s_{\text{assumed}}}$. The first factor characterizes the nature of precision as either equal or variable among models and these represent the two model families in this factor. The second factor considers the dependence of precision parameter on experimental conditions. Specifically, the precision parameter, J (EP) or \bar{J} (VP) can either be condition-independent or condition-dependent. The third factor describes the assumption about $\rho_{s_{\text{assumed}}}$. We have considered following different assumptions about $\rho_{s_{\text{assumed}}}$ in our models: $(0, \frac{1}{3}, \frac{2}{3}, 1)$, $(0, 0, 0, 0)$, $(\alpha, \alpha, \alpha, \alpha)$, $(\alpha, \alpha, \alpha, \beta)$, and $(\alpha, \beta, \gamma, \delta)$.

Therefore, we combine all 20 models in different families based on the underlying factor. We then compare these model families based on Bayesian and Akaike information criteria to find the best fitting model family in the factor. Such

5.1. MODEL FAMILIES

Factor	Model family	Models belonging to the family	Number of models
Precision	EP	EP1 to EP10	10
	VP	VP1 to VP10	10
J or \bar{J}	condition-independent	EP1 to EP5, VP1 to VP5	10
	condition-dependent	EP6 to EP10, VP6 to VP10	10
$\rho_{s_{\text{assumed}}}$	$(0, \frac{1}{3}, \frac{2}{3}, 1)$	EP1, EP6, VP1, VP6	4
	$(0, 0, 0, 0)$	EP2, EP7, VP2, VP7	4
	$(\alpha, \alpha, \alpha, \alpha)$	EP3, EP8, VP3, VP8	4
	$(\alpha, \alpha, \alpha, \beta)$	EP4, EP9, VP4, VP9	4
	$(\alpha, \beta, \gamma, \delta)$	EP5, EP10, VP5, VP10	4

Table 5.1: **List of models based on different factors and model families.** Grouping of models in different model families based on a common assumption about the factor.

a grouping of models in different families or levels helps us in comparing multiple models together which have a common characteristic.

In summary, we group our models in different families based on three different factors. Table 5.1 illustrates a detailed categorization of the models listed in Table 4.1 in different factors and their respective model families.

We compare models across different families based on a common factor and select the model family that is preferable in model comparison. We use both Bayesian information criterion (BIC) and Akaike information criterion (AIC) for model selection. As explained in Section 3.6, these criteria are based on the maximum log-likelihood of a model and penalize the model for its additional free parameters. The penalty term is stronger in the case of BIC as compared to AIC. The model with a lower value of BIC (or AIC) is favorable over a model with a

higher value of BIC (or AIC) in model selection. As an example, we use BIC and AIC to compare equal and variable precision model families and determine which one of them is favorable on the criteria. A favorable model (or model family) has the least value of BIC or AIC.

We usually consider the relative BIC (or AIC) differences between a model and the putative best fitting model. In such a case, a positive relative BIC (or AIC) model difference confirms that the model is worse as compared to the hypothesized best fitting model. The magnitude of relative BIC (or AIC) difference determines the similarity between the predictions of the two models.

In the following sections, we present our model comparison results for different model families belonging to the three different factors. We study both individual and average subject based comparisons. We find that individual differences lead to inconsistent and mixed results in some cases. We first analyze the comparison of models belonging to equal and variable precision families in the first model factor.

5.2 Equal versus variable precision

We first compare our models based on the precision factor and this factor consists of two model families. Each family has 10 models, namely EP1 to EP10 in equal precision and VP1 to VP10 in variable precision model family.

We analyzed the model fits of EP and VP models in Chapter 4. We found that

5.2. EQUAL VERSUS VARIABLE PRECISION

the majority of EP models fitted poorly to the data as compared to their variable precision equivalents. But, some of the EP models generated good predictions for the subjects' responses. For instance, predictions of model EP10 in Figure 4.11 are comparable to those of VP1 (Figure 4.12), while it fits better to the data than model VP2 (Figure 4.13). In such a case, we wish to find the model that best explains the data and is preferable over others in model selection.

Figure 5.1 shows the BIC (top) and AIC (bottom) comparisons of EP models with respect to their VP equivalents. Each vertical bar represents the relative BIC (or AIC) difference of an EP model from its VP equivalent. For example, we consider the relative BIC (or AIC) difference of model EP1 from model VP1 or model EP3 from model VP3. The left panels in Figure 5.1(A) and (B) shows the individual subjects' BIC and AIC comparisons of EP models relative to VP models. We find that all subjects have large positive BIC and AIC relative differences for all EP models. This signifies that variable precision models are preferable for all subjects based on both criteria. Moreover, the difference magnitudes are larger than 500 points for majority of subjects confirming the poor fitting of EP models to the data.

We also consider the average BIC (Figure 5.1(A, right)) and AIC differences (Figure 5.1(B, right)) between EP and VP models. We observe a clear trend of favorable VP models on both criteria. The large positive average relative differences (more than 200 points) provides a strong evidence that variable precision models outperforms their equal precision equivalents when fitting to the data. We thus conclude that subjects use varying precision across stimuli and trials on the target

5.2. EQUAL VERSUS VARIABLE PRECISION

detection experiment.




Figure 5.1: **BIC and AIC model comparisons: equal versus variable precision models. (A) BIC model comparison.** (Left) Relative BIC differences of equal precision models with respect to variable precision models for individual subjects. Each vertical bar represents the BIC difference of an EP model from its equivalent VP model. (Right) Averaged (across subjects) relative BIC differences of equal precision models with respect to variable precision models. **(B) AIC model comparison.** Similar to (A) with AIC relative differences. A positive BIC or AIC relative difference indicates that the EP model is worse as compared to its equivalent VP model. Throughout the chapter, the error bars indicate the unit standard error mean (s.e.m).

5.3 Condition-independent \bar{J} versus condition-dependent \bar{J}

In Section 4.4, we observed that variable precision models have better fits to the data as compared to the equal precision models (Section 4.3). Further, the model comparison results in Figure 5.1 confirm that subjects use varying precision to make decisions on the experiment. These models are strongly preferable in both model selection criteria. Hence, we only consider variable precision models for our further model comparison results.

We now compare variable precision models to determine the dependence of precision parameter on experimental conditions. Specifically, we have two model families in this factor: VP models with mean precision, \bar{J} condition-independent or with \bar{J} condition-dependent. We found in Section 4.4 that condition-dependent models (VP6, VP8, VP9, and VP10) provide better fit to the data, while condition-independent models (VP1 to VP5) fail to explain the responses of subjects on the homogeneous conditions. Based on the model fittings and predictions, models with condition-dependent \bar{J} outperform their equivalents with condition-independence precision assumption. However, the better fits of condition-dependent VP models could be attributed to their additional free parameters (see Table 4.1). Thus, we need to select the best model family among the two.

We compare both these model families of variable precision models: with

5.3. *CONDITION-INDEPENDENT \bar{J} VERSUS CONDITION-DEPENDENT \bar{J}*

condition-independent \bar{J} (models VP1 to VP5) and condition-dependent \bar{J} (models VP6 to VP10) using BIC and AIC in Figure 5.2. The figure shows the BIC and AIC differences of condition-independent models relative to their condition-dependent equivalents. For example, the first vertical bar for the first subject in Figure 5.2(A, left) represents the relative BIC difference of VP1 model with respect to VP6 model (both have a common assumption that $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}}$). Hence, the different color bars represent the following differences: VP1-VP6, VP2-VP7, VP3-VP8, VP4-VP9, and VP5-VP10. Clearly, a positive relative BIC or AIC difference indicates that condition-dependent model is preferable in the comparison.

In Figure 5.2(A, left), we note that the BIC comparisons based on individual subjects are unclear. A total of 5 subjects favor condition-dependent models, while 4 subjects prefer condition-independent models. The other two subjects show mixed preference depending on the different assumptions about ρ_s in different models. Similarly, subject based AIC comparisons in Figure 5.2(B, left) do not give conclusive results for some subjects. Though, majority of subjects clearly favor condition-dependent \bar{J} models with large relative AIC differences. We thus obtain inconsistent results based on the two criteria and they are mainly because of different penalty terms. Bayesian information criterion strongly penalizes a model for its additional free parameters, as compared to Akaike information criterion (Section 3.6). We note that condition-dependent precision models have 3 additional free parameters than condition-independent models (see Table 4.1). Therefore, BIC strongly penalizes these models and we find inconclusive results in this case.

We also analyze the average (across subjects) relative BIC and AIC differences

5.3. CONDITION-INDEPENDENT \bar{J} VERSUS CONDITION-DEPENDENT \bar{J}




Figure 5.2: **BIC and AIC model comparisons: VP models with condition-independent mean precision \bar{J} versus experimental condition-dependent \bar{J} .** (A) **BIC model comparison.** Relative BIC differences of condition-independent \bar{J} VP models with respect to condition-dependent \bar{J} models for each subject (left) and average subject (right). (B) **AIC model comparison.** Similar to (A) with AIC differences.

5.3. CONDITION-INDEPENDENT \bar{J} VERSUS CONDITION-DEPENDENT \bar{J}

of models belonging to the two families in the right panels of Figure 5.2(A) and (B). These comparisons show a clear preference of condition-dependent precision models based on both criteria. Moreover, the difference magnitudes are large. This is because some subjects show a strong preference for condition-dependent model family, hence their individual BIC and AIC differences are relatively large. The large error bars in average model comparisons reflect the individual differences and the inconsistent choice for the model family among subjects.

Therefore, we find that models with condition-dependent \bar{J} are favorable over condition-independent \bar{J} models in average BIC and AIC model selection. This finding is consistent to our observation in Section 4.4 that these models generate better predictions for subjects' psychometric curves. However, it would be difficult to conclude that VP models with condition-dependent \bar{J} represent the best model family since the choice of a model family is different across individual subjects.

In order to compare models based on the assumption about ρ_s , we only consider variable precision models with condition-dependent \bar{J} . In addition to their preference in the average BIC and AIC model comparisons, these are the only models which fitted successfully predicted the behavior of subjects on the experiment.

5.4 Comparison based on $\rho_{S_{assumed}}$

Thus far, we find that variable precision models outperform equal precision models by a huge margin (Figure 5.1) in model selection criteria. Further, we observe that VP models with condition-dependent mean precision are favorable over models with condition-independent \bar{J} (Figure 5.2). Both these results are based on precision factor. We now compare variable precision models based on the assumption about $\rho_{S_{assumed}}$.

The models described in Section 4.2.1 have distinct assumptions about ρ_s and in particular, we consider five different assumptions in our models $\rho_{S_{assumed}} = \rho_{S_{true}}$, $\rho_{S_{assumed}} = (0, 0, 0, 0)$, $\rho_{S_{assumed}} = (\alpha, \alpha, \alpha, \alpha)$, $\rho_{S_{assumed}} = (\alpha, \alpha, \alpha, \beta)$, and $\rho_{S_{assumed}} = (\alpha, \beta, \gamma, \delta)$. Analyzing the model predictions in Sections 4.3 and 4.4, we find that there are multiple models which have equally well fits for the data (see Figures 4.17, 4.19, 4.20, and 4.21). These models are VP6 ($\rho_{S_{assumed}} = \rho_{S_{true}}$), VP8 ($\rho_{S_{assumed}} = (\alpha, \alpha, \alpha, \alpha)$), VP9 ($\rho_{S_{assumed}} = (\alpha, \alpha, \alpha, \beta)$), and VP10 ($\rho_{S_{assumed}} = (\alpha, \beta, \gamma, \delta)$). Thus, there are multiple assumptions of $\rho_{S_{assumed}}$ that equally well predict the behavior of subjects. Therefore, the question arises: which one of these models best explains the data? To answer this question, we compare models VP6 to VP10 using BIC and AIC. We note that all these models have a common assumption about encoding precision and only differ in the hypothesis about $\rho_{S_{assumed}}$.

Figure 5.3 shows the relative BIC (top) and AIC (bottom) differences for individual subjects (left) and average subject (right). The average subject comparison




Figure 5.3: **BIC and AIC model comparisons: VP models with condition-dependent \bar{j} and different assumptions about $\rho_{s_{assumed}}$.** (A) **BIC model comparison.** Relative BIC differences of condition-dependent VP models having different assumptions about $\rho_{s_{assumed}}$ with respect to VP9 model (with $\rho_{s_{assumed}} = (\alpha, \alpha, \alpha, \beta)$) for individual subjects (left) and average subject (right). (B) **AIC model comparison.** Relative AIC differences of condition-dependent VP models with respect to VP10 model (with $\rho_{s_{assumed}} = (\alpha, \beta, \gamma, \delta)$) for individual subjects (left) and average subject (right).

results in Figure 5.3(A, right) and (B, right) show that each criterion selects a different favorable model. Model VP9 has the smallest BIC value in the family, while VP10 is the preferable model based on AIC comparison. But both models only differ by a few points in the selection criteria. The average relative BIC difference between VP9 and VP10 models is 6.91, while it is about 5.46 points based on AIC. This difference is significantly small to distinguish the two models apart and suggests that both models make similar predictions for the subject's responses on the experiment.

Thus, the results are inconsistent based on BIC and AIC comparisons. This is mainly due to the difference between penalty terms of the two model selection criteria. We note that model VP10 has 10 parameters, while VP9 has 8 free parameters (Table 4.1). We recall that each subject performed 3600 trials in the experiment (Section 2.5.3), hence, the penalty term in BIC (Eq. (3.6)) scales each additional parameter in the model by $\log(3600) \approx 8.19$. This amounts to $\log(3600) \times 10 \approx 81.9$ penalty points for a subject in the case of VP10 model. On the other hand, the penalty term is about 65.5 points for VP9 model. Therefore, model comparison based on BIC would prefer the VP9 model in case the maximum log-likelihoods of the two models are comparable. The average subject comparison in the right panel of Figure 5.3(A) illustrates such a preference of VP9 model as compared to the higher dimensional VP10 model.

Similarly, we can analyze the average AIC comparison in Figure 5.3(B, right). The two models only differ by 4 points in AIC penalty term and hence, the most general model VP10 is preferred over others in AIC model selection.

We also find that models VP7 and VP8 having the assumptions of zero and constant correlations perform poorly in both BIC and AIC average model comparisons. The large positive relative differences of these models with respect to the winning models indicate the poor predictions of these models for the data. We observe that model VP6 having the assumption of $\rho_{s_{assumed}} = \rho_{s_{true}}$ also loses to winning models in BIC and AIC average comparisons by 24.4 and 42.3 points, respectively. Therefore, the model selection results confirm that models VP9 and VP10 provide better description of the data over other three condition-dependent VP models. This indicates that subjects make suboptimal inferences on the experiment and use incorrect assumptions about the generative model (described in Figure 2.1) in detecting a target on the task.

However, a large variation in model preference is seen among individuals in the left panels of Figure 5.3(A) and (B). The figure shows the relative BIC differences of other models with respect to VP9 model, while on AIC comparison, the differences are shown with respect to VP10 model. We find that there is no single model that is favored by all subjects. Models VP6, VP9, and VP10 are really close in predictions for majority of subjects and only differ by a few points. This suggest that these models are hard to distinguish on the BIC and AIC measures based on individual differences. Perhaps subjects use different strategies to make decisions on the experiment. However, it is difficult to confirm such a possibility given the limitations of the model comparison techniques we use here.

Therefore, based on the comparison results for models VP7 and VP8 for majority of subjects, we conclude that subjects do take into account the statistical

structure present in visual scenes to make decisions. Further, they also discriminate the different experimental conditions and use different strengths of correlations. However, it is difficult to select the best model between VP9 and VP10 models since they are relatively close in both AIC and BIC comparisons (Figure 5.3). Hence, it is hard to say whether subjects use different strengths of correlations in different experimental conditions (based on VP10 model) or they treat the first three experimental conditions similarly and the homogeneous condition differently as hypothesized by VP9 model.

We further perform a rejection rate analysis to find a single model that best explains the experimental data.

5.5 Rejection rate analysis

Based on the model comparison results in previous sections, it is difficult to select the model that best describes the subjects' behavior on the experiment. Thus, we use rejection rate analysis to determine which model family is preferable in each factor [149]. Table 5.1 lists the three different factors present in our 20 models. Further, each factor has different number of model families that share a common assumption about a parameter in the factor.

We first find the model with least BIC (or AIC) value (among all models) for each subject and it is the most favorable or winning model for that subject. We define the *rejection criterion* as the BIC (or AIC) difference of a model with respect

5.5. REJECTION RATE ANALYSIS

to the winning model of each subject. For each model family in a factor, we compute the number of subjects that reject all its members as a function of the rejection criterion. That is, we count the number of subjects that reject all models belonging to that particular model family for a given criterion. We repeat this process at all values of the BIC or AIC rejection criterion and plot the number of subjects that rejected a particular model family as a function of the criterion in Figure 5.4.

Figure 5.4(A) shows the results for the first model factor, the assumption of encoding precision. Regardless of the rejection criterion, the entire family of equal precision models is rejected for majority of subjects in terms of both BIC (top) and AIC (bottom). Moreover, all subjects reject EP models for a large criterion difference ranging up to 200 points. While, none of the subjects reject the variable precision model family indicating that all individuals have their best models in this model family. This provides a strong evidence that variable precision models are indeed better in describing the behavior of subjects on visual search tasks as found recently in many studies [98, 150, 99, 137, 149].

Figure 5.4(B) shows the rejection analysis for the second factor, the dependence or independence of the encoding precision parameter on experimental conditions. At any given rejection criterion for both BIC (top) and AIC (bottom), most of the subjects select the model family with condition-dependent precision parameter. The number of subjects under a rejection criterion of 0 determines the number of subjects for which the winning model belongs to the model family being considered. For example, in the top panel of Figure 5.4(B), there are 4 subjects that rejected condition-dependent encoding precision models, while 7 rejected the other




Figure 5.4: **Rejection rate curves: comparison of different model families in each model factor.** Each column corresponds to a factor and a color represents a particular model family or level belonging to the factor. The model comparison is based on all 20 models for each factor. **(A) Factor 1: precision.** Number of subjects for whom *all* models belonging to a certain family or level (EP or VP) are rejected as a function of the rejection criterion based on BIC (top) and AIC (bottom) differences. A model is rejected if it has a higher BIC or AIC than that of the winning model for a subject. For example, when both BIC and AIC rejection criteria is 100, all models of the EP family are rejected, while none of the subjects reject VP models. **(B) Factor 2: dependence of precision parameter on experimental conditions.** Similar to (A). **(C) Factor 3: assumption about correlation strength $\rho_{s_{\text{assumed}}}$.** Similar to (A).

family of models. This implies that based on BIC model selection, majority of subjects' responses (63.6%) are best explained by models belonging to the condition-dependent family. On the other hand, 9 subjects have winning model in this family based on AIC differences (bottom panel of Figure 5.4(B)). Further, all subjects prefer model family with condition-dependent precision as the rejection criterion increases. Overall, there is a clear separation between the two model families and subjects more frequently reject the models that assume the encoding precision to be constant across experimental conditions.

Finally, the rejection plots based on the correlation coefficient factor are displayed in Figure 5.4(C). There are five levels in this factor corresponding to the different assumptions about ρ_s . These levels are described in Table 5.1. We observe that the distinction between models is unclear in terms of BIC differences (top). But, the models with the assumption of $\rho_{s_{\text{assumed}}} = (0, 0, 0, 0)$ and $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \alpha)$ are more frequently rejected by subjects at any given rejection criterion for both BIC (top) and AIC (bottom) differences. Next, we note that majority of subjects reject the model with $\rho_{s_{\text{assumed}}} = \rho_{s_{\text{true}}} = (0, \frac{1}{3}, \frac{2}{3}, 1)$ as compared to the models with the assumptions of $\rho_{s_{\text{assumed}}} = (\alpha, \alpha, \alpha, \beta)$ and $\rho_{s_{\text{assumed}}} = (\alpha, \beta, \gamma, \delta)$. This suggests that subjects are suboptimal in inferring the true values of stimulus correlations, ρ_s in the experiment. However, it is still difficult to determine which of the two strategies subjects follow. We see mixed results in terms of both BIC (top) and AIC (bottom) differences. In AIC comparison, the models with $\rho_{s_{\text{assumed}}}$ fitted per condition is less frequently rejected by individuals as compared to the model family that assumes constant correlations across first three experimental

conditions and a different in the case of homogeneous distractors; whereas we see a reverse trend for the choice of model on BIC rejection criterion (top). Therefore, there is a conflict between the choice of the best model based on the two criteria. It may be the case that the mixed results here reflect individual differences: some subjects might have used different correlations across experimental conditions, while some might have treated the conditions with partial correlations identically and the homogeneous condition differently in making their decisions.

These results also agree with our earlier observations about the model predictions for VP9 and VP10 models in Figures 4.20 and 4.21, respectively. Both these variable precision models predict the data equally well, and also their model families show conflicting winning preference in model comparisons (Figure 5.3) and rejection rate analysis (Figure 5.4).

In order to resolve the conflict between the two models, we analyze the maximum-likelihood parameter estimates of the VP9 and VP10 models. The parameter estimates may reflect the cause of similar predictions of the two models.

5.6 Parameter estimates

Parameter estimates for the models VP9 and VP10 are given in Table 5.2. Here \bar{j}_i denotes the estimated mean precision for the gamma distribution of precision in the i^{th} condition. Similarly, ρ_i represents the estimated parameter value of ρ_{assumed} in the i^{th} condition of the experiment. The parameters τ , the scale parameter of

5.6. PARAMETER ESTIMATES

Model	Parameter	Mean \pm SEM	Median
VP9	\bar{J}_1	0.1097 ± 0.0184	0.1036
	\bar{J}_2	0.1077 ± 0.0206	0.1036
	\bar{J}_3	0.1148 ± 0.0227	0.1225
	\bar{J}_4	0.2651 ± 0.0816	0.1447
	ρ_{123}	0.4030 ± 0.0793	0.5
	ρ_4	0.7515 ± 0.1010	0.9333
	τ	0.7433 ± 0.2877	0.3683
	p_T	0.5048 ± 0.0046	0.5
VP10	\bar{J}_1	0.0981 ± 0.0158	0.0877
	\bar{J}_2	0.0937 ± 0.0167	0.0742
	\bar{J}_3	0.1022 ± 0.0185	0.0877
	\bar{J}_4	0.2376 ± 0.0623	0.1447
	ρ_1	0.4333 ± 0.0737	0.5
	ρ_2	0.3939 ± 0.0818	0.4667
	ρ_3	0.4818 ± 0.0896	0.4667
	ρ_4	0.7939 ± 0.0856	0.9333
	τ	0.5691 ± 0.1718	0.3683
	p_T	0.5036 ± 0.0043	0.5

Table 5.2: **Maximum-likelihood parameter estimates of VP9 and VP10 models.** The estimates of the mean precision and scale parameter of the precision distribution are given along with the estimates of correlation strength $\rho_{s_{\text{assumed}}}$ in each experimental condition (where applicable) and of p_T , the observer's prior probability that the target is present.

the gamma distribution of precision, and p_T , the observer's prior probability that the target is present, are estimated as a constant parameter across all conditions in the experiment.

We observe that the mean estimated values of \bar{J}_i are really close for $i = 1, 2, 3$, while the magnitude order of the mean value of \bar{J}_4 differs significantly in both models. This suggests that probably subjects use similar levels of precision in case of experimental conditions with partially correlated stimuli, while they have

higher precision in detecting a target on the homogeneous condition.

The trend in the values of \bar{J} matches our expectation since finding a target among identically oriented distractors would be easier compared to the randomly oriented stimuli. Therefore, an observer would be more precise in the experimental condition with $\rho_s = 1$. Also, we note that subjects use near uniform prior for the target present probability in both models. Further, the maximum-likelihood estimates of the two models suggest that subjects over estimate the low correlation strengths, but perform nearly optimally when distractors are perfectly correlated.

Figure 5.5 shows the behavior of individual parameter estimates for model VP9 in (A) and VP10 in (B). We note that the estimates for both VP9 and VP10 models have a similar trend. The parameters closely follow a similar behavior for majority of subjects and the mean parameter estimates behave nearly identically. The estimates of prior over target presence (first column) and τ (second column) are similar for the two models for most individuals. Also, the values of \bar{J} (third column) looks nearly identical in both models for most of the subjects except for minor deviations. The mean estimated values of \bar{J} show a near constant trend in the first three experimental conditions for the model VP10 (Figure 5.5(B, third column)). Further, the estimates of $\rho_{s_{\text{assumed}}}$ are close in the partial correlation conditions.

A similar trend in parameter estimates of the two models provides an explanation for the equally well model predictions of both models for the data and the conflicting choice of models on BIC and AIC model comparisons.




Figure 5.5: **Individual subject parameter estimates of VP9 and VP10 models.** (A) **Parameter estimates of VP9 model.** Maximum-likelihood estimates of p_T , the observer’s prior probability of target presence (first column) and the scale parameter τ (second column) for each subject. Individual (gray) and average subject (black) estimates of mean precision \bar{J} in different experimental conditions (third column). Estimated values of correlation coefficient ρ_s for each individual (gray) and average subjects (black) as a function of true correlation strength (diagonal) in the experimental conditions. (B) **Parameter estimates of VP10 model.** Similar to (A).

5.7 Conclusions and discussion

Based on the model fittings (Sections 4.3 and 4.4) and model comparisons discussed above, we conclude that subjects use the statistical information of visual scenes to make decisions. That is, they take correlations of the distractors into account in detecting a target on the visual search experiment. However, they are suboptimal in inferring the true correlation strengths that were used to generate the stimuli in different experimental conditions. The model comparison results based on Bayesian information criterion (Figures 5.3(A) and 5.4(C, top)) suggest that perhaps subjects treat the partial correlation conditions identically and find it difficult to distinguish the distinct strengths of weak correlations in the scenes. However, they behave differently when the distractors are homogeneous and show an improved performance in the average results.

While the results based on Akaike information criterion (Figures 5.3(B) and 5.4(C, bottom)) provide evidence that perhaps subjects infer different levels of correlation strengths in different experimental conditions. But they infer incorrect values leading to the suboptimal behavior. Further, we explore that subjects use variable precision in making decisions (Figures 5.1 and 5.4(A)). This precision not only varies over trials and stimuli, but also vary across experimental conditions (Figures 5.2 and 5.4(B)). Thus, we find that distribution of encoding precision depends on the strength of correlation used in an experimental condition.

Our results are inconclusive here about how subjects actually infer the different correlation strengths. This could be because of many underlying reasons. It

may be that the individuals' differences lead to mixed results: subjects probably use different strategies to make decisions on the task. This would be difficult to examine using the measures we used and would probably involve a detailed analysis of each subject's responses. It could also be possible that the models we considered here lack some assumptions about an unknown parameter that could be driving subjects' decisions. For instance, we assumed that subjects learn the parameter value of the standard deviation of distractors in the generative model (Figure 2.1). However, this may not necessarily be true and in that case, we would need models that assume σ_s as a free parameter and the model decision variable would need to be computed by marginalizing over all possible values of this parameter. Further, it could be possible that subjects do not make point estimates of the correlation strengths, instead they use some unknown distribution over the true values. Even worse, another possibility could be that subjects do not use any structural information about in their decisions, but instead use some suboptimal decision strategies, for example, a threshold rule (similar to the one described in Chapter 1, Example 1.4.1).

There could be many other alternatives that are possible here and these may provide a better explanation about subjects' behavior on the experiment. However, all these models are more likely to be further complex and multi-dimensional. The decision variables for these models would be even more sophisticated (compared to Eq. (2.21)) except in the case of threshold models. Therefore, it is a bit uncertain that observers would use even more complex decision variables in their decisions. And if they really do, a mathematical treatment of such models would

5.7. CONCLUSIONS AND DISCUSSION

be highly intricate and may not be possible in many cases. To explore such possibilities, an advance treatment and hierarchal models may be required.

Chapter 6

Measurement correlations in a single target detection task

In many mid- to high-level visual tasks with multiple stimuli, the brain has to make categorical, global judgements. This involves extracting the relevant information from sensory input. Given the nature of a task, the brain needs to process information differently in relevance to the objective of the task. For instance, in a target detection task where the goal is to determine whether a predefined object is present in a visual scene containing multiple objects, the identity of any individual object may not be of direct relevance in making a decision. On the other hand, orientation of each object is equally important in an estimation task where an observer is required to estimate the mean orientation of the presented objects in a visual display. Therefore, the objective of a task could play a crucial role in

guiding our inference process and how our brains extract the meaningful information from sensory input. We explore the relevance of the objective of a task on our decision-making process in some details in Chapter 7.

In addition to a task relevant feature, our judgements also critically depend on the accuracy of our sensory measurements. The sensory information our brain receives is usually uncertain since noise corrupts our measurements, especially when observation time is short and multiple objects are present. The magnitude of noise in the measurements governs the accuracy of our decisions on a task and further complicates the inference process. Thus, measurement noise can considerably affect our performance on the task.

Extensive work has been done both at theoretical and experimental levels to understand the decision processes of the brain. Specifically, several models have been proposed to study the mechanisms by which the brain converts noisy sensory measurements of a set of stimuli to infer the state of the world. For example, how the brain infers a target presence or absence in a scene or how to estimate the mean orientation of a set of stimuli. These models generally consider decision rules that are applied to the measurements. On the other hand, the measurements themselves are usually modeled in a rather conventional fashion. They are often considered as independent (between stimuli) and normally distributed [94, 151, 98, 99]. We have also used these assumptions to model an observer's measurements in Chapter 2. However, both these assumptions can be questioned.

It has been found that neural correlations can extend to distances as long as

4mm in monkey cortex [31, 39]. This indicates that the sensory measurements can be strongly correlated [124, 30]. These correlations must be accounted in the modeling of decision processes. Therefore, we are interested in examining the effects of correlated sensory measurements on our inference process and the accuracy of decisions.

In order to make categorical, global judgements, the brain not only needs to take into account the correlations present in the measurements, but also the statistical structure of the stimuli on a task. For instance, a target could be easily detected in the case of homogeneous distractors as compared to the case when the distractors have random orientations. In such a case, correlations between sensory measurements can further influence decisions. As an example, we consider the situation discussed by Mazzyar et al. [98]. Strong measurement correlations result in more similarity between the measurements and introduce internal structure. This structure could be helpful in presence of identical distractors, as similar measurements corresponding to the distractors can be grouped and a target can be easily detected if present. By contrast, when distractors are independently drawn, there is no external stimulus structure that could be preserved and therefore, we do not necessarily expect strong measurement correlations to be beneficial. This example illustrates that measurement and stimulus correlations should not be considered in isolation.

We expect that correlations between measurements and those between stimuli, will interact to jointly influence the decisions of an observer. In this chapter, we explore the interplay between the measurement and stimulus correlations, and

their effects on the performance of an ideal Bayesian observer in a visual search task. Specifically, we consider a target detection task similar to the one described in Section 2.1, but with the modified assumption of correlated sensory measurements.

We first introduce the model set up for the target detection task and then derive the decision rule that governs the decisions of an optimal Bayesian observer on the task. Our goal here is to understand how external structure together with the structure of the measurements impacts the performance of the ideal Bayesian observer. We examine the performance of the observer in different regimes of various parameters that determine external stimulus and internal measurements' structure.

6.1 Model description

To examine how decisions of an ideal observer are determined by the statistical structure of sensory measurements and stimuli, we consider a single target detection task where the observer is required to detect a vertical stimulus predefined as the target in a set of N stimuli. The framework of the task and mathematical notations are similar to Section 2.1. The binary variable, T , indicates target presence for $T = 1$ and target absence when $T = 0$. The target is present at one of the N possible locations in half of the trials. Stimulus orientations, denoted by $\mathbf{s} = (s_1, s_2, \dots, s_N)$, are the relevant characteristics of the task. We denote the target stimulus orientation by $s_T = 0$ and draw the orientations of the distractors

from multivariate normal distributions given in Eqs. (2.1) and (2.6) when $T = 0$ and $T = 1$, respectively. We note that the external structure of a scene is determined by the number of stimuli, N , the variance of the distractor orientations, σ_s^2 , and the pairwise correlation coefficient, ρ_s . For the purpose of this task, we consider that the number of stimuli and the external noise, σ_s are fixed parameters. We control the amount of structure in a scene by varying the amount of correlation between distractors, ρ_s .

We further assume that an observer makes a noisy *measurement* of each stimulus denoted by x_i . This measurement can be thought of as the maximum likelihood estimate obtained from the activity in a population of neurons with receptive fields including location i . We denote by $\mathbf{x} = (x_1, x_2, \dots, x_N)$, the vector of N measurements. It is commonly assumed that the components of \mathbf{x} are unbiased, independent, and normally distributed [94, 99]

$$x_i | s_i \sim \mathcal{N}(s_i, \sigma_x^2).$$

We consider here the more general situation where the measurements are unbiased, but correlated so that

$$\mathbf{x} | \mathbf{s} \sim \mathcal{N}(\mathbf{s}, \Sigma_x). \quad (6.1)$$

The $N \times N$ covariance matrix, Σ_x is assumed to have a similar structure as covariance matrix Σ_s (Eq. (2.3)) with constant diagonal terms, σ_x^2 and off-diagonal terms, $\rho_x \sigma_x^2$. The assumption that the measurements follow a Gaussian distribution is almost certainly an oversimplification. However, since this distribution is not well characterized, Gaussianity is a reasonable first guess that allows us to describe responses with a minimal number of parameters and leads to analytically

tractable formulations. This assumption is natural when the uncertainty of the measurement is characterized by the variance of each x_i [56, 91]. The structure in the measurements is characterized by the measurement noise, σ_x and the correlation coefficient, ρ_x . Figure 6.1 shows the generative model of the task (A) and the inference process of an optimal Bayesian observer (B). Distributions of the stimuli (A) and the corresponding measurements (B) are illustrated in the case of $N = 2$ stimuli in Figure 6.2.

We now derive the mathematical model for an optimal observer to make decisions on the task. Specifically, we compute the Bayesian decision variable as a function of different statistical parameters that govern the structure in external scenes and the observer's measurements.

6.2 Optimal observer theory

In order to understand how noise correlations in the measurements impact the decisions of an ideal Bayesian observer, we derive an analytical expression for the decision variable following similar computations performed in Section 2.2.

An optimal Bayesian observer makes a decision based on the log posterior ratio given in Eq. (2.9). The observer infers target presence when the decision variable denoted by $d_{\text{NST}}(\mathbf{x})$, is positive, and target absence otherwise. That is,

$$d_{\text{NST}}(\mathbf{x}) = \log \frac{P(T = 1|\mathbf{x})}{P(T = 0|\mathbf{x})} = \underbrace{\log \frac{P(\mathbf{x}|T = 1)}{P(\mathbf{x}|T = 0)}}_{L_{\text{NST}}(\mathbf{x})} + \log \frac{P(T = 1)}{P(T = 0)} > 0. \quad (6.2)$$




Figure 6.1: **Statistical structure of relevant task variables in the optimal-observer model for a (single) target detection task with stimulus and measurement correlations. (A) Generative model.** The binary variable, T indicates target presence for $T = 1$, and absence when $T = 0$. The stimulus orientations, $\mathbf{s} = (s_1, s_2, \dots, s_N)$ are drawn from a multivariate normal distribution with mean vector, $\mathbf{s}_\mathbf{D}$ and covariance matrix, Σ_s . The standard deviation, σ_s and the correlation coefficient, ρ_s of distractor orientations determine the statistical structure of a visual scene. An observer makes measurements, $\mathbf{x} = (x_1, x_2, \dots, x_N)$ of the presented set of stimuli. These measurements are assumed to be unbiased, but correlated, and are drawn from a multivariate normal distribution with a covariance matrix, Σ_x . The correlation coefficient, ρ_x determines the extent of dependence between the sensory measurements. **(B) Inference process.** An ideal Bayesian observer computes the decision variable, $d(\mathbf{x})$ based on the measurements \mathbf{x} to make an estimate, \hat{T} of the true state variable, T . The decision variable, $d(\mathbf{x})$ is the log-posterior ratio between the two possibilities of making a response "target-present" or "target-absent", given the measurements and is given by $\log \frac{p(T=1|\mathbf{x})}{p(T=0|\mathbf{x})}$. The sign of $d(\mathbf{x})$ determines the estimate $\hat{T} = 1$ ($d(\mathbf{x}) > 0$) or $\hat{T} = 0$ ($d(\mathbf{x}) < 0$).




Figure 6.2: **Stimulus and measurement distributions in the single target detection with $N = 2$ stimuli for $\sigma_s = 15^\circ$ and $\sigma_x = 4^\circ$.** (A) **Stimulus distributions for $\rho_s = 0.5$.** Stimulus distributions on target present (left) and target absent (right) trials. On target present trials, one of the two stimuli is uniformly chosen to be the target and hence the distribution is aligned along the target orientation. On target absent trials, the two stimuli have a bivariate normal distribution. (B) **Measurement distributions for $\rho_s = 0.5$.** Measurement distributions for measurement correlations, $\rho_x = 0$ in response to the stimulus distributions in (A) on both target present (left) and target absent (right) trials. On target present trials, the measurement distribution is aligned towards the target orientation (left), while it represents the bivariate normal distribution in target absent cases (right). (C) **Overlap of measurement distributions for $\rho_s = 0.99$.** Measurement distributions on target present (orange) and absent (purple) trials are shown for $\rho_x = 0$ (top), and $\rho_x = 0.95$ (bottom). The overlap between the two measurement distributions, $\mathbf{x}|T = 1$ and $\mathbf{x}|T = 0$ reduces as the strength of measurement correlation, ρ_x increases and the two distributions become more distinguishable. Throughout the chapter, the axes are measured in terms of the standard deviation, σ , which is defined by $\sigma^2 = \sigma_s^2 + \sigma_x^2$.

6.2.1 The log-likelihood ratio

We denote the log-likelihood ratio by $L_{\text{NST}}(\mathbf{x})$ and compute it by marginalizing the information over intermediate variables; the spatial location and the stimulus orientations, \mathbf{s} . We note that it is identical to obtaining Eq. (2.11) in Section 2.2. Therefore,

$$L_{\text{NST}}(\mathbf{x}) = \log \frac{P(\mathbf{x}|T=1)}{P(\mathbf{x}|T=0)} = \log \frac{1}{N} \sum_{j=1}^N \frac{\int P(\mathbf{x}|\mathbf{s})P(\mathbf{s}|T=1)d\mathbf{s}}{\int P(\mathbf{x}|\mathbf{s})P(\mathbf{s}|T=0)d\mathbf{s}}. \quad (6.3)$$

However, we note that the computation of $P(\mathbf{x}|T=1)$ is not simple and straight here than as in Section 2.2. It is because $P(\mathbf{x}|\mathbf{s})$ can no longer be written as a product of one-dimensional normal density functions. Therefore, to evaluate the integral in the case of $T=1$, we construct a new $N \times N$ matrix $\Sigma_{\mathbf{s},j}^\eta$ for $\eta > 0$ as:

$$(\Sigma_{\mathbf{s},j}^\eta)_{k,l} = \begin{cases} \sigma_s^2, & \text{if } k = l \neq j, \\ \rho_s \sigma_s^2, & \text{if } k, l \neq j, \\ \eta, & \text{if } k = l = j, \\ 0, & \text{if } k = j, l \neq j, \text{ or } k \neq j, l = j. \end{cases} \quad (6.4)$$

Here $j \in \{1, 2, \dots, N\}$ represents the spatial location of the target. By denoting $\tilde{\mathbf{s}}_{\mathbf{D}}$ as an N -dimensional vector having all components s_D , but the j^{th} as s_T , we write

$$\mathbf{s}|(T=1, j) \sim \lim_{\eta \rightarrow 0^+} \mathcal{N}(\tilde{\mathbf{s}}_{\mathbf{D}}, \Sigma_{\mathbf{s},j}^\eta). \quad (6.5)$$

We introduced the auxiliary covariance $\Sigma_{\mathbf{s},j}^\eta$ because $\Sigma_{\mathbf{s},j}^0$ is not invertible, as the density is singular in the variable with index j .

In the limit of $\eta \rightarrow 0^+$, we compute the log-likelihood ratio

$$\begin{aligned} L_{\text{NST}}(\mathbf{x}) &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{\int P(\mathbf{x}|\mathbf{s})P(\mathbf{s}|T=1)d\mathbf{s}}{\int P(\mathbf{x}|\mathbf{s})P(\mathbf{s}|T=0)d\mathbf{s}} \right] \\ &= \log \left[\frac{1}{N} \sum_{j=1}^N \frac{\lim_{\eta \rightarrow 0^+} \int f(\mathbf{x}; \mathbf{s}, \Sigma_x) f(\mathbf{s}; \tilde{\mathbf{s}}_D, \Sigma_{s,j}^\eta) d\mathbf{s}}{\int f(\mathbf{x}; \mathbf{s}, \Sigma_x) f(\mathbf{s}; \mathbf{s}_D, \Sigma_s) d\mathbf{s}} \right]. \end{aligned}$$

We use product and integral properties of multivariate normal distributions given in Eqs. (B.3) and (B.4), and define

$$\mathbf{C} = \Sigma_s + \Sigma_x \text{ and } \mathbf{C}_j = \Sigma_{s,j}^0 + \Sigma_x \quad (6.6)$$

to obtain

$$\begin{aligned} L_{\text{NST}}(\mathbf{x}) &= \log \frac{1}{N} \sum_{j=1}^N \frac{f(\mathbf{x}; \tilde{\mathbf{s}}_D, \mathbf{C}_j)}{f(\mathbf{x}; \mathbf{s}_D, \mathbf{C})} \\ &= \log \left[\frac{1}{N} \sqrt{\frac{|\mathbf{C}|}{|\mathbf{C}_j|}} \sum_{j=1}^N \exp \left(-\frac{1}{2} (\mathbf{x} - \tilde{\mathbf{s}}_D)^T \mathbf{C}_j^{-1} (\mathbf{x} - \tilde{\mathbf{s}}_D) \right. \right. \\ &\quad \left. \left. + \frac{1}{2} (\mathbf{x} - \mathbf{s}_D)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{s}_D) \right) \right]. \quad (6.7) \end{aligned}$$

The determinant of the matrix \mathbf{C}_j does not depend on the spatial location parameter j since all matrices of type \mathbf{C}_j can be obtained from each other by permuting appropriate rows and columns. Moreover, the distributions of the Gaussian variables, $\mathbf{x}|(T=0)$ and $\mathbf{x}|(T=1)$ are described by their respective covariances given by matrices \mathbf{C} and \mathbf{C}_j .

6.2.1.1 Determinants and inverses of matrices \mathbf{C} and \mathbf{C}_j

We further simplify Eq. (6.7) by evaluating the determinants and inverses of the covariance matrices \mathbf{C} and \mathbf{C}_j , when $\rho_s, \rho_x \neq 1$; otherwise these matrices are singular. We first decompose these matrices as:

$$\mathbf{C} = \mathbf{D} + (\rho_s \sigma_s^2 + \rho_x \sigma_x^2) \mathbf{J} \text{ and } \mathbf{C}_j = \mathbf{A}_j + \mathbf{U}_j \mathbf{E} \mathbf{V}_j,$$

where \mathbf{D} is a diagonal matrix with constant diagonal entries, $\sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x)$, \mathbf{J} is a matrix of ones, \mathbf{E} is a 2×2 identity matrix,

$$\text{and } (\mathbf{A}_j)_{k,l} = \begin{cases} \eta + \sigma_x^2(1 - \rho_x), & \text{if } k = l = j, \\ \sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x), & \text{if } k = l \neq j, \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, the columns of $N \times 2$ matrix \mathbf{U}_j and rows of $2 \times N$ matrix \mathbf{V}_j are given by

$$\mathbf{U}_{j(k,1)} = \begin{cases} \sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x), & \text{if } k \neq j, \\ \sigma_x^2(1 - \rho_x), & \text{if } k = j, \end{cases} \text{ and } \mathbf{U}_{j(k,2)} = \sigma_x^2(1 - \rho_x) \quad \forall k,$$

$$\mathbf{V}_{j(1,k)} = \begin{cases} 1, & \text{if } k \neq j, \\ 0, & \text{if } k = j, \end{cases} \text{ and } \mathbf{V}_{j(2,k)} = \begin{cases} 0, & \text{if } k \neq j, \\ 1, & \text{if } k = j. \end{cases}$$

Additionally, we define the following variables

$$v = \frac{1}{\sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x)}, \quad \tilde{v} = \frac{1}{\sigma_x^2(1 - \rho_x)}, \quad a = \rho_s \sigma_s^2 + \rho_x \sigma_x^2, \quad (6.8)$$

$$V = Nv, \quad V_{\setminus j} = (N - 1)v, \quad \beta = \frac{a}{1 + aV}, \quad q = a + \rho_s \sigma_s^2 \rho_x \sigma_x^2 \tilde{v}, \quad (6.9)$$

$$r = 1 + \rho_s \sigma_s^2 V_{\setminus j}, \quad \text{and } \gamma = 1 + aV_{\setminus j} + \rho_x \sigma_x^2 \tilde{v} r. \quad (6.10)$$

6.2. OPTIMAL OBSERVER THEORY

We apply the generalized Matrix Determinant Lemma (Lemma 2 in Appendix B.3) to obtain the following determinants:

$$|\mathbf{C}| = \frac{1 + aV}{v^N}, \text{ and } |\mathbf{C}_j| = \frac{\gamma}{v^{N-1} \tilde{v}}.$$

We use the Woodbury matrix identity (Theorem 2 in Appendix B.3) to compute the following inverses of both covariance matrices

$$(\mathbf{C}^{-1})_{k,l} = \begin{cases} v - \beta v^2 & \text{if } k = l, \\ -\beta v^2 & \text{if } k \neq l, \end{cases}$$

$$\text{and } (\mathbf{C}_j^{-1})_{k,l} = \begin{cases} \tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 r}{\gamma} & \text{if } k = l = j, \\ v - \frac{v^2 q}{\gamma} & \text{if } k = l \neq j, \\ -\frac{v^2 q}{\gamma} & \text{if } k \neq l, \text{ and } k, l \neq j, \\ -\frac{\rho_x \sigma_x^2 v \tilde{v}}{\gamma} & \text{if } k = j, l \neq j, \text{ or } l = j, k \neq j. \end{cases}$$

We now substitute the above obtained determinants and inverses of \mathbf{C} and \mathbf{C}_j to compute the following expressions:

$$(\mathbf{x} - \mathbf{s}_D)^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{s}_D) = (v - \beta v^2) \sum_{k=1}^N (x_k - s_D)^2 - \beta v^2 \sum_{k \neq l}^N (x_k - s_D)(x_l - s_D)$$

$$(\mathbf{x} - \mathbf{s}_{\tilde{D}})^T \mathbf{C}_j^{-1} (\mathbf{x} - \mathbf{s}_{\tilde{D}}) = \left(\tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 r}{\gamma} \right) (x_j - s_T)^2 + v \sum_{k \neq j}^N (x_k - s_D)^2$$

$$- \frac{2\rho_x \sigma_x^2 v \tilde{v}}{\gamma} (x_j - s_T) \sum_{k \neq j}^N (x_k - s_D) - \frac{v^2 q}{\gamma} \sum_{k, l \neq j}^N (x_k - s_D)(x_l - s_D).$$

Substituting the above obtained expressions in Eq. (6.7) gives us the following simplified equation for the log-likelihood ratio on the single target detection task

with correlated measurements

$$\begin{aligned}
 L_{\text{NST}}(\mathbf{x}) = \log & \left[\frac{1}{N} \sqrt{\frac{\tilde{v}(1+aV)}{v\gamma}} \sum_{j=1}^N \exp \left(-\frac{1}{2} \left(\tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 r}{\gamma} \right) (x_j - s_T)^2 \right. \right. \\
 & - \frac{1}{2} \left(\beta v^2 - v \right) (x_j - s_D)^2 - \left(\beta v^2 - \frac{\rho_x \sigma_x^2 v \tilde{v}}{\gamma} \right) (x_j - s_D) \sum_{k \neq j}^N (x_k - s_D) \\
 & \left. \left. - \frac{1}{2} \left(\beta v^2 - \frac{v^2 q}{\gamma} \right) \sum_{k,l \neq j}^N (x_k - s_D)(x_l - s_D) \right) \right]. \tag{6.11}
 \end{aligned}$$

6.2.2 Bayesian decision variable

The decision variable, $d_{\text{NST}}(\mathbf{x})$ defined in Eq. (6.2) characterizes the decision-making process of an ideal Bayesian observer on the task. It also quantifies the impact of different (stimulus and measurement structure) parameters on the decisions of the optimal observer.

Since a target is present with uniform probability on each trial of the task, the ideal observer uses a uniform prior on T . Therefore, the decision variable to report "target present" on the task is equivalent to the log-likelihood ratio obtained in Eq. (6.11),

$$\begin{aligned}
 d_{\text{NST}}(\mathbf{x}) = \log & \left[\frac{1}{N} \sqrt{\frac{\tilde{v}(1+aV)}{v\gamma}} \sum_{j=1}^N \exp \left(\underbrace{-\frac{1}{2} \left(\tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 r}{\gamma} \right) (x_j - s_T)^2}_I \right. \right. \\
 & \left. \left. - \underbrace{\frac{1}{2} (\beta v^2 - v) (x_j - s_D)^2}_{II} - \underbrace{\left(\beta v^2 - \frac{\rho_x \sigma_x^2 v \tilde{v}}{\gamma} \right) (x_j - s_D) \sum_{k \neq j}^N (x_k - s_D)}_{III} \right. \right. \\
 & \left. \left. - \underbrace{\frac{1}{2} \left(\beta v^2 - \frac{v^2 q}{\gamma} \right) \sum_{k,l \neq j}^N (x_k - s_D)(x_l - s_D)}_{IV} \right) \right]. \tag{6.12}
 \end{aligned}$$

The above equation gives the decision variable in terms of the model parameters and the measurements, \mathbf{x} . The total number of stimuli, N , the variability, and co-variability between the distractors' orientation determined respectively by σ_s^2 and ρ_s , regulate the structure of visual stimuli. The variability, σ_x^2 and correlation strength, ρ_x , describe the structure of the observer's measurements.

6.2.3 Interpretation of the decision variable

Eq. (6.12) defines a nonlinear decision boundary in the space of measurements, \mathbf{x} . Although the expression is explicit, the decision variable depends in a complicated way on the different parameters that describe the structure of the stimulus and the response. The variables v and \tilde{v} represent scaled inverse variances corresponding to distractor and target stimulus, while V and $V_{\setminus j}$ denote the population sums of v in target present and absent cases, respectively. The parameters β, r, q ,

and γ (Eq. (6.8)) are defined in terms of $\sigma_s^2, \rho_s, \sigma_x^2$, and ρ_x , and it is difficult to quantify their explicit impact on decisions and performance of an observer.

However, we note that the formulation of Eq. (6.12) is similar to the decision variable derived in Eq. (2.21), except for the different scaling parameters. Thus, Eq. (6.12) can be interpreted in a similar way as in Section 2.2.2. Specifically, we interpret each exponent term of Eq. (6.12) as an evidence towards the j^{th} stimulus being a target: (I) if the j^{th} measurement is close to the target orientation, this term increases the likelihood of the j^{th} stimulus being a target; (II) on the other hand, if the j^{th} measurement is similar to the mean distractor orientation, this term decreases such a likelihood; (III) the third term compares the sample mean of distractor measurements with the j^{th} measurement; if the term is large, it is less likely that the j^{th} stimulus is the target; and (IV) the fourth term can be rewritten as the sample covariance of distractor measurements; if this term is large it is more likely that the j^{th} stimulus is the target. Thus, each term is relevant in making a decision whether the j^{th} stimulus is a target or not. The coefficient of each exponent term contain information about various parameters that govern the stimulus and measurement structure. The influence of these parameters such as $\sigma_s^2, \rho_s, \sigma_x^2$, and ρ_x on different terms of the decision variable, $d_{\text{NST}}(\mathbf{x})$ is difficult to understand, since their dependence is much more intricate.

We next consider a simple and particular case of the decision variable in Eq. (6.12). Specifically, we evaluate the equation in the absence of measurement correlations.

6.2.4 Uncorrelated measurements, $\rho_x = 0$

We note that when sensory measurements are assumed to be uncorrelated, *i.e.*, for $\rho_x = 0$, the mathematical model of the task reduces to the one described in Section 2.1. The variables defined in Eq. (6.8) reduces to the ones defined in Chapter 2 (Eqs. (2.15) and (2.19)) on substituting $\rho_x = 0$,

$$a = \rho_s \sigma_s^2, v = \tilde{w}, \tilde{v} = w, V = \tilde{W}, \beta = \alpha, \frac{q}{\gamma} = \alpha_{\setminus j}.$$

Subsequently, the decision variable, $d_{\text{NST}}(\mathbf{x})$, given in Eq. (6.12), reduces to the variable, $d_{\text{ST}}(\mathbf{x})$, derived in Eq. (2.21). Thus, the target detection task described in Chapter 2 is a particular case of the task discussed here having a generalized assumption about measurement correlations.

We now consider how variations in the external structure parameters together with the parameters governing the measurements' structure impact the performance of an ideal Bayesian observer. External structure is characterized by the number of stimuli and the pairwise correlations between them, while the structure in the measurements is specified by the covariance matrix, Σ_x . In general, we cannot assume that these parameters can be varied independently. For example, neural mechanisms that impact σ_x^2 almost certainly impact ρ_x [31, 36, 123]. However, the dependence between these parameters is not fully characterized and we therefore explore a range of possible parameter values below.

6.3 Analysis and Results

Our goal is to describe how stimulus structure, along with the structure of the corresponding measurements, affect the decisions of an optimal observer in a target detection task. We examine how performance changes as stimulus and measurement correlations are varied.

We first note that the stimulus variable, \mathbf{s} , follows two different distributions depending on whether $T = 0$ or $T = 1$. In general, measurement noise increases the overlap between these distributions. The higher the overlap between the two distributions, the more difficult the decision. However, structured noise in the measurements can reduce such overlap (Figure 6.2(C)). Therefore, performance of an ideal observer depends not only on the level, but also on the structure of measurement noise [6, 4].

We examine Eq. (6.12) numerically in different regimes of stimulus and measurement statistical parameters. Specifically, we consider the following two regimes determined by the dominance of either external noise, σ_s or measurement noise, σ_x :

(a) weak measurement noise, $\sigma_x^2 \ll \sigma_s^2$,

(b) strong measurement noise, $\sigma_x^2 = \sigma_s^2$.

In the following sections, we elaborate how noise correlations in the measurements affect the decision variable, $d_{\text{NST}}(\mathbf{x})$ (Eq. (6.12)) and subsequently the performance of an ideal Bayesian observer. We obtain analytical expressions for the

decision variable in these cases and provide supporting numerical results along with the best possible intuitive interpretation.

6.3.1 Weak measurement noise, $\sigma_x^2 \ll \sigma_s^2$

The external structure in a scene is determined by the noise level σ_s and also by the pairwise correlation coefficient, ρ_s . For a fixed noise level σ_s , the structure is introduced in the visual stimuli by varying the amount of correlations among pairs of distractor orientations. A strong external structure is introduced when $\rho_s = 1$, as all distractors are identical and the target is an odd-ball stimulus. In the cases of $\rho_s < 1$, the external structure is weaker, hence detecting a target may not necessarily be easier. Therefore, we individually treat the cases of weak and strong external structure in the regime of weak noise in the observer's measurement.

We denote $\epsilon = \frac{\sigma_x^2}{\sigma_s^2}$ and expand different terms of Eq. (6.12) about $\epsilon \ll 1$. We only consider terms with larger contribution, *i.e.*, terms of the orders of $\mathcal{O}(\frac{1}{\epsilon})$ to obtain approximations for the decision variable, $d_{\text{NST}}(\mathbf{x})$ in different regimes of parameters.

6.3.1.1 Strong external structure, $\rho_s = 1$

In the case of $\rho_s = 1$ and $\epsilon \ll 1$, the coefficients of the different exponent terms in Eq. (6.12) reduces to the following simplified expressions:

$$\begin{aligned}
 I : \tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 r}{\gamma} &\approx \frac{1}{\sigma_x^2} \\
 II : \beta v^2 - v &\approx \frac{-(N-1)}{N \sigma_x^2 (1 - \rho_x)} \\
 III : \beta v^2 - \frac{\rho_x \sigma_x^2 v \tilde{v}}{\gamma} &\approx \frac{1}{N \sigma_x^2 (1 - \rho_x)} + \mathcal{O}(1) \\
 IV : \beta v^2 - \frac{v^2 q}{\gamma} &\approx -\frac{1}{N(N-1) \sigma_x^2 (1 - \rho_x)},
 \end{aligned}$$

and the leading determinant coefficient becomes

$$\sqrt{\frac{\tilde{v}(1 + aV)}{v \gamma}} \approx \sqrt{\frac{N(1 - \rho_x)}{N - 1}}.$$

Therefore, we obtain the following approximation of the Bayesian decision variable, $d_{\text{NST}}(\mathbf{x})$ in this case

$$\begin{aligned}
 d_{\text{NST}}(\mathbf{x}) \approx \log &\left[\frac{1}{N} \sqrt{\frac{N(1 - \rho_x)}{N - 1}} \sum_{j=1}^N \exp \left(-\frac{1}{2N \sigma_x^2 (1 - \rho_x)} \left((1 - N \rho_x)(x_j - s_T)^2 \right. \right. \right. \\
 &- (N - 1)(x_j - s_D)^2 + 2(x_j - s_D) \sum_{k \neq j}^N (x_k - s_D) \\
 &\left. \left. \left. - \frac{1}{N - 1} \sum_{k, l \neq j}^N (x_k - s_D)(x_l - s_D) \right) \right) \right]. \tag{6.13}
 \end{aligned}$$

Furthermore, in the particular case of $s_T = s_D = 0$, the above equation reduces to a much simpler expression that is easier to interpret

$$d_{\text{NST}}(\mathbf{x}) \approx \log \left[\frac{1}{N} \sqrt{\frac{N(1 - \rho_x)}{N - 1}} \sum_{j=1}^N \exp \left(\underbrace{-\frac{1}{2N\sigma_x^2(1 - \rho_x)} \left((1 - N\rho_x)x_j^2 + 2x_j \sum_{k \neq j}^N x_k - \frac{1}{N - 1} \left(\sum_{k \neq j}^N x_k \right)^2 \right)}_{E_j} \right) \right]. \quad (6.14)$$

Clearly, the above expression indicates that an ideal observer uses the strength of measurement correlations, ρ_x in a decision. Specifically, in the limit of perfect noise correlations, *i.e.*, $\rho_x \approx 1$, the exponential term in Eq. (6.14) is approximately

$$E_j \approx \left(\frac{N - 1}{2N\sigma_x^2(1 - \rho_x)} \left(x_j - \frac{1}{N - 1} \sum_{k \neq j}^N x_k \right)^2 \right). \quad (6.15)$$

Therefore, the optimal observer simply subtracts the mean of the $N - 1$ measurements of putative distractors from that of the putative target. In the case of perfect noise correlations, the measurements of the distractor stimuli are identical. Thus, on "target absent" trials, the term E_j 's in Eq. (6.15) are zero and therefore, $d_{\text{NST}}(\mathbf{x}) \rightarrow -\infty$. While on a target present trial, E_j is positive and the exponential diverges to infinity as $\rho_x \rightarrow 1$. The prefactor in Eq. (6.14) approaches zero, however, the divergence of exponential is stronger. Therefore, $d_{\text{NST}}(\mathbf{x}) \rightarrow \infty$ and an ideal observer performs perfectly in this case.

We confirm the above analytical observation using numerical simulations in the case of $N = 4$ stimuli. We vary the strength of stimulus and measurement

noise correlations (ρ_s and ρ_x , respectively) and examine the trend in the performance of an optimal Bayesian observer shown in Figure 6.3(A). For simulation purposes, we used $\sigma_s = 15^\circ$ and $\sigma_x = 4^\circ$ to obtain the weak measurement noise ($\sigma_x^2 \ll \sigma_s^2$) regime.

We observe in Figure 6.3(A) that the performance of an ideal observer in the case of weak measurement noise is nearly independent of ρ_s and ρ_x for weak external structure ($\rho_s < 1$). However, it depends strongly on ρ_x when the external structure is relatively stronger, *i.e.*, $\rho_s \approx 1$. In particular, any amount of measurement correlations strongly drives the performance in such a case (seen in Figure 6.3(A) and (B)). Moreover, we note that perfect performance is observed at $\rho_s = \rho_x = 1$ (Figure 6.3(B)). Hence, in the presence of strong external structure (all distractors having identical orientations), strong measurement correlations enhance the performance of an optimal Bayesian observer.

The increased performance with increasing external correlations, ρ_s accords with intuition that strong external structure makes detecting a target easier. However, *measurement* structure can play an equally important role. In the presence of strong external structure (homogeneous distractors), measurement correlations can significantly improve performance as seen in Figure 6.3(B).

Perfect performance is achieved at $\rho_s = \rho_x = 1$ and we can understand it intuitively. In this case, measurements x_i of the stimuli are obtained by adding the *same* realization of a random variable (noise) to each stimulus characteristic, s_i . In target absent trials, all measurements are hence identical. If the target is present, the vector of measurements contain a single outlier and an ideal observer




Figure 6.3: Performance of an ideal Bayesian observer in the presence of measurement and stimulus correlations on a single target detection task when the measurement noise is weak as compared to the external noise, $\sigma_x^2 \ll \sigma_s^2$ ($\sigma_s = 15^\circ$, $\sigma_x = 4^\circ$). (A) Variation in performance with measurement and stimulus correlations. Performance of an optimal observer as a function of stimulus correlations, ρ_s and measurement correlations, ρ_x on a task with $N = 4$ stimuli. (B) Change in performance with measurement correlations. Proportion of correct responses as a function of measurement correlations, ρ_x in the case of weak external structure, $\rho_s = 0.5$ (left) and strong external structure, $\rho_s = 1$ (right) for $N = 4$ stimuli in the task. (C) Decision boundary and measurement distributions for $N = 2$ stimuli and $\rho_s = 0.5$. Decision boundary, $d_{\text{NST}}(\mathbf{x}) = 0$ (black) and measurement distributions corresponding to correct inferences (orange) and incorrect inferences (purple) on target present (left) and target absent (right) cases. The green curves (lines on left, ellipses on right) represent 2 units standard deviation of the stimulus distribution.

can thus distinguish the two cases perfectly.

6.3.1.2 Weak external structure, $\rho_s < 1$

We also observe that measurement correlations have little effect on performance when external structure is weaker, *i.e.*, when $\rho_s < 1$ (Figure. 6.3(A)). To intuitively understand such a behavior, we again consider the case of $\rho_x \approx 1$ so that measurements are obtained by equal rotation of each stimulus orientation. An ideal observer can use the structure of measurement noise in making a decision. However, the observer cannot use knowledge about a *particular realization* of noise. In other words, the observer can use the fact that the measurements are obtained by rotating the stimulus by approximately the same angle, but not the exact angle of the rotation. If there is only weak external structure, it is now difficult to tell whether one of the stimuli is an outlier. An ideal observer must therefore infer whether a target is present from the individual measurements of stimuli. Hence measurement correlations provide little help in the absence of external structure.

These observations are also reflected in the structure of the decision boundaries ($d_{\text{NST}}(\mathbf{x}) = 0$) and distributions of the measurements illustrated in Figure 6.3(C). In the target present (Figure 6.3(C, left)) and absent trials (right), the distribution of measurements is determined predominantly by variability in the stimulus. Measurement correlations have little effect on these distributions. As a result, the decision boundary also changes little with an increase in ρ_x . This is in contrast to the case when $\rho_s \approx 1$, where internal variability can be important in increasing the overlap between the distributions of $\mathbf{x}|T = 1$ and $\mathbf{x}|T = 0$ (see Figure 6.2(C)).

We confirm this intuition about the role of measurement correlations by approximating the decision variable in this case and expanding the coefficients of exponent terms in Eq. (6.12) as follows

$$\begin{aligned}
 I : \bar{v} - \frac{\rho_x \sigma_x^2 \bar{v}^2 r}{\gamma} &\approx \frac{1}{\sigma_x^2} \\
 II : v - \beta v^2 &\approx \frac{1 + (N - 2)\rho_s}{\sigma_s^2(1 - \rho_s)(1 + (N - 1)\rho_s)} = \mathcal{O}(1) \\
 III : \beta v^2 - \frac{\rho_x \sigma_x^2 v \bar{v}}{\gamma} &\approx \mathcal{O}(1) \\
 IV : \beta v^2 - \frac{v^2 q}{\gamma} &\approx \mathcal{O}(1).
 \end{aligned}$$

We further approximate the leading determinant coefficient in Eq. (6.12) as

$$\sqrt{\frac{\bar{v}(1 + aV)}{v \gamma}} \approx \sqrt{\frac{\sigma_s^2(1 - \rho_s)(1 + (N - 1)\rho_s)}{\sigma_x^2(1 + (N - 2)\rho_s)}}$$

and combine the above terms to obtain the following approximation of the decision variable when $\sigma_x^2 \ll \sigma_s^2$ and $\rho_s < 1$,

$$d_{\text{NST}}(\mathbf{x}) \approx \log \left[\frac{1}{N} \sqrt{\frac{\sigma_s^2(1 - \rho_s)(1 + (N - 1)\rho_s)}{\sigma_x^2(1 + (N - 2)\rho_s)}} \sum_{j=1}^N \exp \left(-\frac{(x_j - s_T)^2}{2\sigma_x^2} \right) \right]. \quad (6.16)$$

This approximation is relatively simpler and easier to understand. We easily see that the strength of measurement correlations, ρ_x does not affect the decisions of an ideal observer at highest order in (σ_x^2/σ_s^2) . Measurement correlations only weakly impact the decision and hence performance. Additionally, the terms in the exponent are $x_j^2/2\sigma_x^2$. Hence, the orientation at each location is considered separately and weighted by the precision of measurements, $1/\sigma_x^2$. The ideal observer hence makes a decision primarily based on the evidence from each stimulus separately in this case.

In sum, external structure of a scene implies that stimulus distributions $\mathbf{s}|T = 1$ and $\mathbf{s}|T = 0$ are concentrated on low-dimensional subspaces. Even small noise can increase the overlap between the distributions of measurements $\mathbf{x}|T = 1$ and $\mathbf{x}|T = 0$. However, noise in the measurement and external structure can conspire to decrease the overlap between the distribution of measurements. This happened in the present case when ρ_s and ρ_x are both close to 1 (Figure 6.3(B)).

In the absence of external structure, the signal distributions $\mathbf{s}|T = 1$ and $\mathbf{s}|T = 0$ are not concentrated along low-dimensional manifolds. Here external variability, σ_s^2 always dominates and low intensity measurement variability, σ_x^2 , has little effect on the performance of an ideal observer.

6.3.2 Strong measurement noise, $\sigma_x^2 = \sigma_s^2$

With a single target, measurement correlations have little impact on the performance of an ideal observer, unless external structure is strong. Although, external variability is typically expected to be much stronger than measurement variability in most situations [14], we next consider the case when the two sources of variability are comparable, $\sigma_x^2 = \sigma_s^2$. The case of $\sigma_x^2 \gg \sigma_s^2$ is insignificant to consider, since the increased measurement noise leads to poor performance and the observer essentially makes a guess on each trial.

Increasing measurement noise, σ_x , trivially leads to a decrease in performance. However, in the limit of perfect stimulus and measurement correlations, *i.e.*, $\rho_s, \rho_x \approx$

1, an ideal observer still performs perfectly. But, it is not clear how the performance is affected at intermediate values of these correlations. An analytical approximation of Eq. (6.12) is difficult to obtain in this particular case. We thus explore this case numerically and provide an intuitive reasoning for the observed behavior. For simulation purposes, we consider the case of $N = 4$ stimuli and $\sigma_s = \sigma_x = 15^\circ$. Figure 6.4(A) shows the performance of an optimal observer with varying strengths of stimulus, ρ_s and measurement correlations, ρ_x .

On comparing Figure 6.4(A), (B) with Figure 6.3(A), (B), we observe that measurement correlations have a strikingly different effect here than in the case of weak measurement noise. Even with no external structure, $\rho_s = 0$, performance increases slightly (approximately 5-6%) with an increase in ρ_x . Surprisingly, for intermediate values of external correlations, $\rho_s = 0.5$, measurement correlations have a negative impact on performance. The reason for this unexpected behavior is not clear, as Eq. (6.12) is difficult to analyze in this case. However, it is obvious that structure present in the observer's measurements play a much greater role in decision making when noise is very large, unlike the weak noise case where we see little change in structure when ρ_x is varied (Figure 6.3(A)).

We see this fact reflected in the distributions of the responses and the decision boundary in Figure 6.4(C), which are both strongly affected by measurement correlations ρ_x . Moreover, we see a similarity between this decision boundary and the one shown in Figure 6.3(C), in that the mid-section is elongated along the diagonal to capture more of the response distribution with an increase in measurement correlations. This transformation makes intuitive sense, considering the fact that




Figure 6.4: **Performance of an ideal Bayesian observer in the presence of measurement and stimulus correlations on a single target detection task when the measurement noise is comparable to the external noise, $\sigma_x^2 = \sigma_s^2$ ($\sigma_s = \sigma_x = 15^\circ$).** (A) **Variation in performance with measurement and stimulus correlations.** Performance of an optimal observer as a function of stimulus correlations, ρ_s and measurement correlations, ρ_x on a task with $N = 4$ stimuli. (B) **Change in performance with measurement correlations.** Proportion of correct responses as a function of measurement correlations, ρ_x for $\rho_s = 0$ (left) and $\rho_s = 0.5$ (right) on the task with $N = 4$ stimuli. (C) **Decision boundary and measurement distributions for $N = 2$ stimuli and $\rho_s = 0.5$.** Decision boundary, $d_{NST}(\mathbf{x}) = 0$ (black) and measurement distributions corresponding to correct inferences (orange) and incorrect inferences (purple) on target present (left) and target absent (right) cases. The green curves (lines on left, ellipses on right) represent 2 units standard deviation of the stimulus distribution.

with increased measurement correlations, noise simply moves the stimulus elements roughly along the diagonal and therefore the decision boundary changes accordingly.

6.4 Summary

The performance of an ideal observer on a detection task with a single target is greatly influenced by the joint interaction of the statistical structure of the stimuli, and the correlations between the observer's measurements. The decisions are largely affected by how different parameters that control the stimulus and measurements structure are varied. Specifically, we find that the performance of the observer is mostly unaffected at all levels of measurements correlations when the stimulus structure is not strong and the measurement noise is weaker than the external noise. However, the observer always make correct decisions when distractors are homogeneous ($\rho_s = 1$) and measurements are perfectly correlated. This is because identical distractors induces a strong statistical structure in a scene, and strongly correlated measurements preserve this structure and help in making correct decisions. In the case of strong measurement noise which is comparable to the external noise ($\sigma_x^2 = \sigma_s^2$), the change in performance for weak external structure is little unpredictable and largely depends on the interaction between the stimulus and measurement correlations. While, the observer still performs perfectly when both distractors and measurements are maximally correlated.

Measurement correlations in a multiple target detection task

In a detection task with a single target, the impact of measurement correlations on the performance of an ideal Bayesian observer strongly depends on the relation between the statistics governing the external structure and the parameters associated with the observer's measurements. Measurement correlations have varying influence on the decisions of the ideal observer in different parameter regimes, as seen in Chapter 6.

Specifically, the interaction between these correlations is beneficial when the amount of structure in visual scenes is sufficiently strong and pronounced (Figure 6.3(A)). In the case of weak measurement noise, measurement correlations increase performance (Figure 6.3(B)) when external structure is strong. However, in the case of strong measurement noise, the decisions of the ideal observer are

also affected when external structure is weak (Figure 6.4(A)). Therefore, the impact of measurement correlations on the accuracy of decisions heavily depends on the amount of structure present in an external visual scene.

The external statistical structure in a target detection task is determined by various factors: the number of stimuli, the number of targets, the mean orientation of the target and distractors, the external noise level in the distractors, and the pairwise correlation coefficient between distractor orientations. By keeping all other factors fixed, the amount of structure in a scene can be introduced by varying the amount of correlations between distractor orientations; namely from heterogeneous to homogeneous distractors (as seen in Chapter 2). However, we can also vary any other parameter that drives the framework of a visual scene and then study how measurement correlations could possibly interact with the structured input to affect the performance of an ideal observer.

One possibility is to increase the number of stimuli in the task. The increased number of stimuli and hence the distractors, could possibly help an ideal observer in the case of homogeneous distractors; however, in general we would expect a lower performance when the distractors are not sufficiently aligned, as in the case of heterogeneous distractors. This is because the observer would be required to integrate more information from additional stimuli and their respective locations in such a case.

Rather, we consider an alternative to introduce more structure in the external scene. We allow the possibility of more than one target in a visual display. Since

multiple targets can be present in a display, we refer to such a task as *multiple target detection task*. For a given number of stimuli in the task, increasing the number of targets certainly enhances the structure in the visual scene and can potentially benefit the decisions of an ideal observer. However, the impact of measurement correlations is unknown in the case of such structured inputs. We thus explore how an optimal observer would behave in presence of stimulus and measurement correlations on a multiple target detection task. In particular, we analyze the responses of the ideal observer in different regimes of parameters determining the structure of the observer's measurements and the visual scene.

Furthermore, we inspect whether measurement correlations always reinforce the decisions of an observer or could they possibly hurt them? Could a different objective, based on a same set of stimuli, lead to different behavioral decision? This is an important question and we have briefly discussed it in the introduction of Chapter 6. We examine this question in a *discrimination task*. The task for the observer is to determine whether the mean orientation of the displayed stimuli is to the left or right of the vertical. We deliberately design the discrimination task with similar structural characteristics as that of the multiple target detection task. Thus, the two tasks have similar structural displays but differ in the objective for the observer. The aim is to study how measurement correlations affect the performance of the optimal observer on both tasks and how their influence vary with the objective of the task.

We begin this chapter with model description of the multiple target detection task, followed by the mathematical derivation of the optimal-observer model. We

further analyze the influence of measurement correlations on the distributions of the observer's responses and the resulting performance. We conclude by comparing the effects of measurement correlations in the case of a discrimination task to that of the multiple target detection task. We find that measurement correlations have an adverse effect on the performance in the discrimination task.

7.1 Model description

The model set up for a multiple target detection task is very similar to the single target detection task described in Sections 2.1 and 6.1. We present a set of N stimuli to the observer and the observer needs to infer whether a target stimulus is present in the visual display. A target is a stimulus with a particular orientation, denoted by s_T . For simplicity, we assume $s_T = 0^\circ$ and measure stimulus orientations relative to that of a target. On half of the trials, one or multiple targets are present among a set of distractors. All n targets have identical orientation. A distractor stimulus has a non-target orientation. Again, we denote target presence by $T = 1$ and absence by $T = 0$.

When $T = 0$, there are no targets and stimulus orientations are drawn from a multivariate normal distribution described in Eq. (2.1) with mean $\mathbf{s}_D = \mathbf{0}_N = (0, 0, \dots, 0)$ and covariance matrix, Σ_s , so that

$$\mathbf{s}|T = 0 \sim \mathcal{N}(\mathbf{0}_N, \Sigma_s).$$

We again assume Σ_s with constant diagonal terms, σ_s^2 and off-diagonal terms,

7.1. MODEL DESCRIPTION

$\rho_s \sigma_s^2$, where the pairwise correlation coefficient, ρ_s determines the dependence between orientations of the distractors.

If $T = 1$, then $n \geq 1$ targets are present. To place the targets, we choose n out of the N possible locations with equal probability. The set of the n target locations is denoted by L . There are $\binom{N}{n}$ possible choices for this set and we denote the collection of all possible choices of sets of target locations by \mathcal{L} . We also write $M = \binom{N}{n}$ as the cardinality of the set \mathcal{L} .

The orientations of the remaining $N - n$ distractors are chosen from a multivariate normal distribution with mean $\mathbf{0}_{N-n}$ and covariance matrix $\Sigma_{\mathbf{s}_{\setminus L}}$ of dimension $(N - n) \times (N - n)$. Once the locations of the targets are chosen, let $\mathbf{s}_L = (s_{i_1}, s_{i_2}, \dots, s_{i_n})$, $i_l \in L$ denote the orientations of the target stimuli and $\mathbf{s}_{\setminus L} = (s_{j_1}, s_{j_2}, \dots, s_{j_{N-n}})$, $j_l \notin L$ those of the distractors. We can therefore write

$$\mathbf{s}_L | T = 1 \sim \sum_{i \in L} \delta(s_i) \quad \text{and} \quad \mathbf{s}_{\setminus L} | T = 1 \sim \mathcal{N}(\mathbf{0}_{N-n}, \Sigma_{\mathbf{s}_{\setminus L}}). \quad (7.1)$$

For $\eta > 0$, we introduce the following auxiliary covariance,

$$(\Sigma_{\mathbf{s}, L}^\eta)_{i,j} = \begin{cases} (\Sigma_{\mathbf{s}_{\setminus L}})_{i,j}, & \text{if } i, j \notin L, \\ \eta, & \text{if } i = j \in L, \\ 0, & \text{if } i \in L, \text{ or } j \in L, \text{ and } i \neq j, \end{cases} \quad (7.2)$$

and write,

$$\mathbf{s} | (T = 1, L) \sim \lim_{\eta \rightarrow 0^+} \mathcal{N}(\mathbf{0}_N, \Sigma_{\mathbf{s}, L}^\eta).$$

We note that the matrix $\Sigma_{\mathbf{s}, L}^\eta$ reduces to $\Sigma_{\mathbf{s}, j}^\eta$ defined in Eq. (6.4) in the case of a detection task with a single target, $n = 1$.

We further assume that the observer's measurements, $\mathbf{x} = (x_1, x_2, \dots, x_N)$ are correlated so that it follows the multivariate normal distribution described in Eq. (6.1) as

$$\mathbf{x}|\mathbf{s} \sim \mathcal{N}(\mathbf{s}, \Sigma_{\mathbf{x}}).$$

Additionally, we assume that the covariance matrix, $\Sigma_{\mathbf{x}}$, has a similar structure as $\Sigma_{\mathbf{s}}$, with constant diagonal terms, σ_x^2 and off-diagonal terms, $\rho_x \sigma_x^2$. Since the task variables are similar to the target detection task specified in Section 6.1, the generative model remains the same as shown in Figure 6.1.

7.2 Optimal observer theory

We now derive the mathematical model for the decisions of an optimal Bayesian observer on the task. In Sections 2.2 and 6.2, we described that the ideal observer makes a decision based on the sign of the log-posterior ratio. Here we denote the Bayesian decision variable by $d_{\text{NMT}}(\mathbf{x})$,

$$d_{\text{NMT}}(\mathbf{x}) = \log \frac{p(T = 1|\mathbf{x})}{p(T = 0|\mathbf{x})}. \quad (7.3)$$

Since the optimal observer uses a uniform prior over T , we essentially compute the log-likelihood ratio

$$d_{\text{NMT}}(\mathbf{x}) = \log \frac{p(\mathbf{x}|T = 1)}{p(\mathbf{x}|T = 0)}.$$

The optimal observer needs to marginalize over the spatial location vector and the stimulus, \mathbf{s} , to compute the distribution of the measurements given the target

presence variable, *i.e.*, $p(\mathbf{x}|T)$. Thus, we have

$$d_{\text{NMT}}(\mathbf{x}) = \log \frac{\int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|T=1)d\mathbf{s}}{\int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|T=0)d\mathbf{s}}.$$

We note that

$$p(\mathbf{s}|T=1) = \sum_{L \in \mathcal{L}} p(\mathbf{s}|T=1, L)p(L) = \frac{1}{M} \sum_{L \in \mathcal{L}} p(\mathbf{s}|T=1, L).$$

Therefore,

$$\begin{aligned} p(\mathbf{x}|T=1) &= \frac{1}{M} \int p(\mathbf{x}|\mathbf{s}) \sum_{L \in \mathcal{L}} p(\mathbf{s}|T=1, L) d\mathbf{s} \\ &= \frac{1}{M} \sum_{L \in \mathcal{L}} \int p(\mathbf{x}|\mathbf{s}) p(\mathbf{s}|T=1, L) d\mathbf{s} \\ &= \frac{1}{M} \lim_{\eta \rightarrow 0^+} \sum_{L \in \mathcal{L}} \int f(\mathbf{x}; \mathbf{s}, \Sigma_{\mathbf{x}}) f(\mathbf{s}; \mathbf{0}_N, \Sigma_{\mathbf{s}, L}^{\eta}) d\mathbf{s} \\ &= \frac{1}{M} \sum_{L \in \mathcal{L}} f(\mathbf{x}; \mathbf{0}_N, \mathbf{C}_L). \end{aligned}$$

Here we have defined $\mathbf{C}_L = \Sigma_{\mathbf{x}} + \Sigma_{\mathbf{s}, K}^0$. The above equation is obtained by applying the general rule of product and integral for multivariate normal distributions described in Appendix B.1.2. Similarly, we define $\mathbf{C} = \Sigma_{\mathbf{x}} + \Sigma_{\mathbf{s}}$ to obtain

$$\begin{aligned} p(\mathbf{x}|T=0) &= \int p(\mathbf{x}|\mathbf{s}) p(\mathbf{s}|T=0) d\mathbf{s} \\ &= \int f(\mathbf{x}; \mathbf{s}, \Sigma_{\mathbf{x}}) f(\mathbf{s}; \mathbf{0}_N, \Sigma_{\mathbf{s}}) d\mathbf{s} \\ &= f(\mathbf{x}; \mathbf{0}_N, \mathbf{C}). \end{aligned}$$

Taking the required ratio of $P(\mathbf{x}|T=1)$ and $P(\mathbf{x}|T=0)$ gives us

$$\begin{aligned} d_{\text{NMT}}(\mathbf{x}) &= \log \frac{1}{M} \sum_{L \in \mathcal{L}} \frac{f(\mathbf{x}; \mathbf{0}_N, \mathbf{C}_L)}{f(\mathbf{x}; \mathbf{0}_N, \mathbf{C})} \\ &= \log \left[\frac{1}{M} \sqrt{\frac{|\mathbf{C}|}{|\mathbf{C}_L|}} \sum_{L \in \mathcal{L}} \exp \left(-\frac{1}{2} \mathbf{x}^T (\mathbf{C}_L^{-1} - \mathbf{C}^{-1}) \mathbf{x} \right) \right]. \quad (7.4) \end{aligned}$$

We note that the determinant of the matrix \mathbf{C}_L does not depend on the set L , since all matrices of this form can be obtained from each other by permuting appropriate rows and columns.

We further simplify the above equation by computing the determinant and inverse of matrix \mathbf{C}_L . The determinant and inverse of \mathbf{C} has already been computed in Section 6.2.1.1. We note that the covariance matrix \mathbf{C}_L has the following form

$$(\mathbf{C}_L)_{i,j} = \begin{cases} \eta + \sigma_x^2, & \text{if } i = j \in L, \\ \sigma_s^2 + \sigma_x^2, & \text{if } i = j \notin L, \\ \rho_x \sigma_x^2, & \text{if either } i, \text{ or } j \in L, \text{ and } i \neq j, \\ \rho_s \sigma_s^2 + \rho_x \sigma_x^2, & \text{if } i, j \notin L, \text{ and } i \neq j. \end{cases}$$

We follow the inverse and determinant computations performed in Section 6.2.1.1.

In particular, we decompose $\mathbf{C}_L = \mathbf{A}_L + \mathbf{U}_L \mathbf{E} \mathbf{V}_L$. The different matrix components in this decomposition are defined in a similar manner as seen in Section 6.2.1.1.

Specifically, \mathbf{E} is a 2×2 identity matrix,

$$(\mathbf{A}_L)_{k,l} = \begin{cases} \eta + \sigma_x^2(1 - \rho_x), & \text{if } k = l \in L, \\ \sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x), & \text{if } k = l \notin L, \\ 0, & \text{otherwise,} \end{cases}$$

and the columns of $N \times 2$ matrix \mathbf{U}_L and rows of $2 \times N$ matrix \mathbf{V}_L are given by

$$\mathbf{U}_{L(k,1)} = \begin{cases} \sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x), & \text{if } k \notin L, \\ \sigma_x^2(1 - \rho_x), & \text{if } k \in L, \end{cases} \quad \text{and } \mathbf{U}_{L(k,2)} = \sigma_x^2(1 - \rho_x), \quad \forall k,$$

$$\mathbf{V}_{L(1,k)} = \begin{cases} 1, & \text{if } k \notin L, \\ 0, & \text{if } k \in L, \end{cases} \quad \text{and } \mathbf{V}_{L(2,k)} = \begin{cases} 0, & \text{if } k \notin L, \\ 1, & \text{if } k \in L. \end{cases}$$

We also recall the definition of the following quantities from Eq. (6.8)

$$v = \frac{1}{\sigma_s^2(1 - \rho_s) + \sigma_x^2(1 - \rho_x)}, \quad \tilde{v} = \frac{1}{\sigma_x^2(1 - \rho_x)},$$

$$V = Nv, \quad a = \rho_s \sigma_s^2 + \rho_x \sigma_x^2, \quad \beta = \frac{a}{1 + aV},$$

and further define

$$V_{\setminus L} = (N - n)v, \quad \tilde{q} = a + n\rho_s \sigma_s^2 \rho_x \sigma_x^2 \tilde{v}, \quad (7.5)$$

$$\tilde{r} = 1 + \rho_s \sigma_s^2 V_{\setminus L}, \quad \text{and } \tilde{\gamma} = 1 + aV_{\setminus L} + n\rho_x \sigma_x^2 \tilde{v} \tilde{r}. \quad (7.6)$$

Applying the Generalized Matrix Determinant Lemma and the Woodbury formula (Appendix B.3), we obtain the following determinant and inverse of the matrix \mathbf{C}_L :

$$|\mathbf{C}_L| = \frac{\tilde{\gamma}}{\tilde{v}^n v^{N-n}},$$

and

$$(\mathbf{C}_L^{-1})_{i,j} = \begin{cases} \tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{r}}{\tilde{\gamma}}, & \text{if } i = j \in L, \\ v - \frac{v^2 \tilde{q}}{\tilde{\gamma}}, & \text{if } i = j \notin L, \\ -\frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{r}}{\tilde{\gamma}}, & \text{if } i, j \in L, \text{ and } i \neq j, \\ -\frac{v^2 \tilde{q}}{\tilde{\gamma}}, & \text{if } i, j \notin L, \text{ and } i \neq j, \\ -\frac{\rho_x \sigma_x^2 v \tilde{v}}{\tilde{\gamma}}, & \text{if } i \neq j, i \in L, j \notin L, \text{ or } i \notin L, j \in L. \end{cases}$$

As required in Eq. (7.4), we simplify

$$\sqrt{\frac{|\mathbf{C}|}{|\mathbf{C}_L|}} = \sqrt{\frac{1 + aV}{\tilde{\gamma}} \left(\frac{\tilde{v}}{v}\right)^n}$$

and compute

$$\begin{aligned} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} &= (v - \beta v^2) \sum_{i=1}^N x_i^2 - \beta v^2 \sum_{i \neq j}^N x_i x_j, \\ \mathbf{x}^T \mathbf{C}_L^{-1} \mathbf{x} &= \left(\tilde{v} - \frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{r}}{\tilde{\gamma}} \right) \sum_{i \in L} x_i^2 + \left(v - \frac{v^2 \tilde{q}}{\tilde{\gamma}} \right) \sum_{i \notin L} x_i^2 \\ &\quad - \frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{r}}{\tilde{\gamma}} \sum_{\substack{i, j \in L \\ i \neq j}} x_i x_j - \frac{2\rho_x \sigma_x^2 v \tilde{v}}{\tilde{\gamma}} \sum_{\substack{i \in L \\ j \notin L}} x_i x_j - \frac{v^2 \tilde{q}}{\tilde{\gamma}} \sum_{\substack{i, j \notin L \\ i \neq j}} x_i x_j. \end{aligned}$$

We substitute the above expressions in Eq. (7.4) to obtain the following expression for the optimal Bayesian decision variable on the task:

$$\begin{aligned}
 d_{\text{NMT}}(\mathbf{x}) = \log & \left[\frac{1}{M} \sqrt{\frac{1+aV}{\tilde{\gamma}}} \left(\frac{\tilde{v}}{v}\right)^n \sum_{L \in \mathcal{L}} \exp \left(-\frac{1}{2} \underbrace{\sigma_s^2 (1 - \rho_s) v \tilde{v} \sum_{i \in L} x_i^2}_I \right. \right. \\
 & - \frac{1}{2} \underbrace{\left(\beta v^2 - \frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{\gamma}}{\tilde{\gamma}} \right) \sum_{i,j \in L} x_i x_j}_{II} - \underbrace{\left(\beta v^2 - \frac{\rho_x \sigma_x^2 v \tilde{v}}{\tilde{\gamma}} \right) \sum_{i \in L, j \notin L} x_i x_j}_{III} \\
 & \left. \left. - \frac{1}{2} \underbrace{\left(\beta v^2 - \frac{v^2 \tilde{q}}{\tilde{\gamma}} \right) \sum_{i,j \notin L} x_i x_j}_{IV} \right) \right]. \tag{7.7}
 \end{aligned}$$

The above equation describes the decision strategy of an optimal observer on the multiple target detection task. If $d_{\text{NMT}}(\mathbf{x}) > 0$, the observer responds the target is present, and target absent otherwise. Eq. (7.7) shows the intricate dependence of the decision variable, $d_{\text{NMT}}(\mathbf{x})$ on several parameters governing the structure of the task such as the total number of stimuli N , number of targets n , variance σ_s^2 , and correlation ρ_s of the distractors' orientations and the parameters σ_x^2 and ρ_x determining the structure of the measurements.

7.2.1 Interpretation of the decision variable

The decision variable derived in Eq. (7.7) is a generalized form of the decision variable computed in Eq. (6.12). Eq. (7.7) is further complex and depends in a complicated way on different parameters that describe the structure of the stimulus and the response. However, we can interpret it in a similar manner as in

Section 6.2.3. The outer sum is over all putative sets of targets $L \in \mathcal{L}$. The summands correspond to the evidence that the set L contains targets. Each term in the exponent has an intuitive interpretation: (I) for the putative set of targets, L , the first term represents the sample second moment of the potential target measurements; if this term is large, the measurements are more likely to be away from zero and hence it is less likely that the set L consists of targets; (II) the second term can be written in terms of sample covariance between putative targets and a smaller value of this term corresponds to an increased chance of set L being a set of targets; (III) the third term compares the sample means of measurements in the putative target set to those outside of it, a small value of this term provides evidence that the set L contains targets; and (IV) the last term can be rewritten in terms of sample covariance of the measurements corresponding to the putative distractor set, *i.e.*, stimuli outside of set L , if these measurements are correlated, a large covariance would imply that this is a set of distractors, hence it is more likely that the set L consists of targets. Again, it is difficult to provide a precise interpretation of the different prefactors involved in Eq. (7.7), since they have more complicated dependence on various parameters in the generative model of the task.

We now analyze the impact of different statistical parameters that govern the structure of a scene and the measurements on the performance of an ideal Bayesian observer in the multiple target detection task. Specifically, we aim to understand how the relationship between the parameters that determine the external structure of the visual display, namely, σ_x^2 and ρ_x , and the measurement parameters, σ_x^2

and ρ_x , affect the decisions of the optimal observer. Since no explicit dependence is known among these parameters, we therefore examine a range of possibilities here.

7.3 Analysis and results

Here we analyze the decision variable, $d_{\text{NMT}}(\mathbf{x})$ (Eq. (7.7)), in the regime of weak measurement noise, that is, when $\sigma_x^2 \ll \sigma_s^2$. In the case of strong measurement noise, $\sigma_x^2 \gg \sigma_s^2$, the noise dominates over other statistical parameters and the observer only makes a guess about target presence. Therefore, we limit our analysis in the regime of weak measurement noise and examine the impact of measurement correlations on the performance of a Bayes-optimal observer.

As seen in Section 6.3.1, we let $\epsilon = \frac{\sigma_x^2}{\sigma_s^2}$ and approximate Eq. (7.7) in the limit of $\epsilon \rightarrow 0$. We split our analysis in the case of weak and strong external structure, that is, when $\rho_s < 1$, and $\rho_s \approx 1$, respectively.

7.3.1 Weak external structure, $\rho_s < 1$

In the absence of structured visual scenes, the coefficients of the exponent terms in Eq. (7.7) reduce to the following expressions in the limit of $\epsilon \rightarrow 0$,

$$\begin{aligned}
 I : \sigma_s^2(1 - \rho_s)v\tilde{v} &= \frac{1}{\sigma_x^2(1 - \rho_x)} + \mathcal{O}(1) \\
 II : \beta v^2 - \frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{\gamma}}{\tilde{\gamma}} &= -\frac{\rho_x}{\sigma_x^2(1 - \rho_x)(1 + (n - 1)\rho_x)} + \mathcal{O}(1) \\
 III : \beta v^2 - \frac{\rho_x \sigma_x^2 v\tilde{v}}{\tilde{\gamma}} &\approx \mathcal{O}(1) \\
 IV : \beta v^2 - \frac{v^2 \tilde{q}}{\tilde{\gamma}} &\approx \mathcal{O}(1).
 \end{aligned}$$

Also, the leading determinant term approximates to

$$\sqrt{\frac{1 + aV}{\tilde{\gamma}} \left(\frac{\tilde{v}}{v}\right)^n} \approx \sqrt{\frac{(1 - \rho_x)(1 + (N - 1)\rho_s)}{(1 + (N - n - 1)\rho_s)(1 + (n - 1)\rho_x)} \left(\frac{\sigma_s^2(1 - \rho_s)}{\sigma_x^2(1 - \rho_x)}\right)^n}.$$

Therefore, we obtain the following approximation for the decision variable in the parameter regime of $\sigma_x^2 \ll \sigma_s^2$ and $\rho_s < 1$:

$$\begin{aligned}
 d_{\text{NMT}}(\mathbf{x}) &\approx \log \left[\frac{1}{M} \sqrt{\frac{(1 - \rho_x)(1 + (N - 1)\rho_s)}{(1 + (N - n - 1)\rho_s)(1 + (n - 1)\rho_x)} \left(\frac{\sigma_s^2(1 - \rho_s)}{\sigma_x^2(1 - \rho_x)}\right)^n} \right. \\
 &\quad \left. \times \sum_{L \in \mathcal{L}} \exp \left(-\frac{1}{2\sigma_x^2(1 - \rho_x)} \left(\sum_{i \in L} x_i^2 - \frac{\rho_x}{1 + (n - 1)\rho_x} \sum_{i, j \in L} x_i x_j \right) \right) \right]. \quad (7.8)
 \end{aligned}$$

Special case: $n = 1$ We note that in the case of a single target, set L has only one element and $\mathcal{L} = \{1, 2, \dots, N\}$. Thus, when $s_T = s_D = 0^\circ$, Eq. (7.8) reduces to the Eq. (6.16), which has been examined in detail in Chapter 6.

We observe that Eq. (7.8) explicitly depends on ρ_x , unlike the case of a single target in the same regime (Section 6.3.1.2). Also, these measurement correlations

clearly influence the performance of an ideal observer, as seen in Figure 7.1(A). The simulations were performed with $N = 4$ stimuli and $n = 3$ targets in the task. The external noise and measurement noise were fixed at $\sigma_s = 15^\circ$ and $\sigma_x = 4^\circ$. Since the n target stimuli have identical orientation, they are perfectly correlated. This introduces a relation between a subset of the stimuli on half of the trials, even when $\rho_s = 0$.

In general, for a fixed ρ_s , we observe in Figure 7.1(A) that the performance increases with increasing ρ_x . This is because the stimuli have a strong structure as all the targets are identically oriented and thus the distribution $\mathbf{s}|T = 1$ lies on a low-dimensional subspace. Therefore, we expect that measurement correlations can have a significant impact in the presence of such structure. This is in addition to the trivial increase in performance expected simply from having more targets.

As seen in Figure 7.1(A) and (B), the ideal observer takes into account measurement correlations for all values of ρ_s , even when measurement noise is low. This lends to an increase in performance, as ρ_x is increased beyond some critical value for any ρ_s . This is different from the single target case (Eq. (6.16)), where the decision variable was independent of measurement correlations when $\rho_s < 1$.

When $\rho_s < 1$, performance gradually increases with ρ_x . The target stimuli are a perfectly correlated subset of the stimulus set. The measurements of these stimuli are identical when measurement correlations are perfect. In such a case, an ideal observer performs perfectly by detecting whether any n of the N measurements x_i are equal.




Figure 7.1: **Performance of an ideal Bayesian observer in the presence of measurement and stimulus correlations on a multiple target detection task when the measurement noise is weak as compared to the external noise, $\sigma_x^2 \ll \sigma_s^2$ ($\sigma_s = 15^\circ$, $\sigma_x = 4^\circ$).** (A) **Variation in performance with measurement and stimulus correlations.** Performance of an optimal observer as a function of stimulus correlations, ρ_s and measurement correlations, ρ_x on a task with $N = 4$ stimuli and $n = 3$ targets. (B) **Change in performance with measurement correlations.** Proportion of correct responses as a function of measurement correlations, ρ_x for $\rho_s = 0$ (left), and $\rho_s = 0.5$ (right) on the task with $N = 4$ stimuli and $n = 3$ targets. (C) **Decision boundary and measurement distributions for $N = n = 2$ and $\rho_s = 0.5$.** Decision boundary, $d_{\text{NST}}(\mathbf{x}) = 0$ (black) and measurement distributions corresponding to correct inferences (orange) and incorrect inferences (purple) on target present (left) and target absent (right) cases. The green ellipses represent 2 units standard deviation of the stimulus distribution and the axes are measured in terms of the standard deviation, σ , which is defined by $\sigma^2 = \sigma_s^2 + \sigma_x^2$.

We now analyze the impact of varying amounts of measurement correlations on the performance of the ideal observer in presence of weak external structure, $\rho_s < 1$, as seen in Figure 7.1(A).

7.3.1.1 No measurement correlations, $\rho_x = 0$

In the absence of measurement correlations, *i.e.*, $\rho_x = 0$, a decision is solely based on the sample second moment *about the target orientation* of the measurements of the n stimuli in the putative target set

$$d_{\text{NMT}}(\mathbf{x}) \approx \log \left[\frac{1}{M} \sqrt{\frac{(1 + (N - 1)\rho_s)}{(1 + (N - n - 1)\rho_s)(1 + (n - 1)\rho_s)}} \left(\frac{\sigma_s^2(1 - \rho_s)}{\sigma_x^2} \right)^n \right. \\ \left. \times \sum_{L \in \mathcal{L}} \exp \left(-\frac{1}{2\sigma_x^2} \sum_{i \in L} x_i^2 \right) \right]. \quad (7.9)$$

Since we have assumed the target orientation to be 0° , this value does not appear explicitly in the centered second moment in Eq. (7.9). If i is a target, the variance of the measurement x_i is σ_x^2 and if i is a distractor, the variance is $\sigma_s^2 + \sigma_x^2$. Therefore, if a subset, L of stimuli contains targets, the centered sample second moment will be smaller than if it does not. Hence, in the absence of measurement correlations, the ideal observer does not take into account the relation between the measurements and instead compares the stimulus orientations in the putative set of targets to the known target orientation.

7.3.1.2 Perfect measurement correlations, $\rho_x \rightarrow 1$

On the other hand, when the measurements are structured, *i.e.*, in the case of $\rho_x > 0$, the ideal observer computes the second moment about a point between the target orientation and the sample mean, $\frac{1}{n} \sum_{i \in L} x_i$. This can be seen by rewriting Eq. (7.8) as

$$d_{\text{NMT}}(\mathbf{x}) \approx \log \left[\frac{1}{M} \sqrt{\frac{(1 - \rho_x)(1 + (N - 1)\rho_s)}{(1 + (N - n - 1)\rho_s)(1 + (n - 1)\rho_s)} \left(\frac{\sigma_s^2(1 - \rho_s)}{\sigma_x^2(1 - \rho_x)} \right)^n} \right. \\ \left. \times \sum_{L \in \mathcal{L}} \exp \left(- \frac{n}{2\sigma_x^2(1 - \rho_x)} \left(\frac{1}{n} \sum_{i \in L} x_i^2 - \frac{n\rho_x}{1 + (n - 1)\rho_x} \left(\frac{1}{n} \sum_{i \in L} x_i \right)^2 \right) \right) \right]. \quad (7.10)$$

Furthermore, in the limit of $\rho_x \rightarrow 1$, the underlined term in Eq. (7.10) approaches the sample variance – the sample second moment centered at the sample mean,

$$\frac{1}{n} \sum_{i \in L} x_i^2 - \frac{n\rho_x}{1 + (n - 1)\rho_x} \left(\frac{1}{n} \sum_{i \in L} x_i \right)^2 \rightarrow E_L = \frac{1}{n} \sum_{i \in L} x_i^2 - \left(\frac{1}{n} \sum_{i \in L} x_i \right)^2. \quad (7.11)$$

In the case of strong measurement correlations, the target stimuli in set L have approximately equal measurements. Therefore, the optimal observer computes the sample second moment centered around the sample mean (the sample variance).

We first consider the case of target present trials, $T = 1$. We note that E_L in Eq. (7.11) approaches zero for a set of targets, L_T and therefore, the exponential in Eq. (7.10) becomes unity. In the event that L is a set not consisting of all targets, $E_L > 0$ and

$$\sum_{L \in \mathcal{L}} \exp \left(- \frac{nE_L}{2\sigma_x^2(1 - \rho_x)} \right) = 1 + \sum_{L \in \mathcal{L} \setminus L_T} \exp \left(- \frac{nE_L}{2\sigma_x^2(1 - \rho_x)} \right) \rightarrow 1.$$

The prefactor in Eq. (7.10) diverges. However, since $\lim_{x \rightarrow \infty} \exp(-x)x^n = 0$ for any finite n , the product of the prefactor and the exponential terms (not equal to 1) still converge to zero. Therefore, on target present trials, we obtain $d_{\text{NMT}}(\mathbf{x}) \rightarrow \infty$ as $\rho_x \rightarrow 1$.

On target absent trials, $T = 0$, the expression E_L will be greater than zero for all sets L . Thus, $d_{\text{NMT}}(\mathbf{x}) \rightarrow -\infty$ in the limit of $\rho_x \rightarrow 1$. Hence, regardless of external correlations ρ_s , the ideal observer performs perfectly when $\rho_x = 1$.

7.3.1.3 Intermediate measurement correlations, $0 < \rho_x < 1$

When $0 < \rho_x < 1$, the ideal observer makes a decision by adopting an intermediate strategy and computes a second moment about a point between the target orientation and the sample mean (Eq. (7.8)). Moreover, the weight on the sample mean increases with the number of targets in set L , since the prefactor $n\rho_x/(1 + (n - 1)\rho_x)$ grows with n for fixed ρ_x .

We also observe in Figure 7.1(C) that measurement correlations significantly impact the distribution of measurements. With $N = n = 2$, the distribution approaches the diagonal ($x_1 = x_2$) as $\rho_x \rightarrow 1$. This decreases the overlap between the target present and absent distributions.

Therefore, with multiple targets, there is always structure in the stimulus set, and $P(\mathbf{s}|T = 1)$ is always concentrated on a low-dimensional subspaces. In this case, the structure present in the observer's measurements can again decrease the

overlap between the distribution of measurements and significantly impact decisions and hence, the performance.

7.3.2 Strong external structure, $\rho_s = 1$

In Figure 7.1(A), we also observe that the performance steadily increases with ρ_x when $\rho_s = 1$ and a perfect performance is achieved at $\rho_s = \rho_x = 1$. This behavior is quite similar to the single target case seen in Figure 6.3(A). We thus expect a similar approximation of the decision variable, $d_{\text{NMT}}(\mathbf{x})$ in this case.

In the case of homogeneous distractors, $\rho_s = 1$, we obtain the following reduced expressions for the different exponent terms in Eq. (7.7):

$$\begin{aligned}
 I : \sigma_s^2(1 - \rho_s)v\tilde{v} &= 0 \\
 II : \beta v^2 - \frac{\rho_x \sigma_x^2 \tilde{v}^2 \tilde{\gamma}}{\tilde{\gamma}} &= \frac{1 - (N - n + 1)\rho_x}{N(1 + (n - 1)\rho_x)\sigma_x^2(1 - \rho_x)} + \mathcal{O}(1) \\
 III : \beta v^2 - \frac{\rho_x \sigma_x^2 v \tilde{v}}{\tilde{\gamma}} &= \frac{1}{N\sigma_x^2(1 - \rho_x)} + \mathcal{O}(1) \\
 IV : \beta v^2 - \frac{v^2 \tilde{q}}{\tilde{\gamma}} &= \frac{-n - N(n - 1)\rho_x}{N(N - n)\sigma_x^2(1 - \rho_x)} + \mathcal{O}(1).
 \end{aligned}$$

The determinant prefactor becomes

$$\sqrt{\frac{1 + aV}{\tilde{\gamma}} \left(\frac{\tilde{v}}{v}\right)^n} \approx \sqrt{\frac{N(1 - \rho_x)}{(N - n)(1 + (n - 1)\rho_x)}}.$$

Combining the above expressions gives us the following approximated decision

variable:

$$d_{\text{NMT}}(\mathbf{x}) \approx \log \left[\frac{1}{M} \sqrt{\frac{N(1-\rho_x)}{(N-n)(1+(n-1)\rho_x)}} \sum_{L \in \mathcal{L}} \exp \left(-\frac{1}{2} \frac{1}{N\sigma_x^2(1-\rho_x)} \right. \right. \\ \left. \left. \underbrace{\left(\frac{1-(N-n+1)\rho_x}{(1+(n-1)\rho_x)} \sum_{i,j \in L} x_i x_j + 2 \sum_{i \in L, j \notin L} x_i x_j - \frac{n+N(n-1)\rho_x}{(N-n)} \sum_{i,j \notin L} x_i x_j \right)}_{F_L} \right) \right]. \quad (7.12)$$

We note that the above equation reduces to Eq. (6.14) studied in Chapter 6 in the case of a single target, $n = 1$.

It is evident from Eq. (7.12) that ρ_x impacts the decisions of an optimal observer when external structure is predominant. When $\rho_x = 1$, the expression F_L in Eq. (7.12) becomes

$$F_L = \frac{n-N}{n} \left(\sum_{i \in L} x_i \right)^2 + 2 \sum_{i \in L, j \notin L} x_i x_j - \frac{n+N(n-1)}{(N-n)} \left(\sum_{i, j \notin L} x_i \right)^2.$$

Therefore, the ideal observer computes the decision in this case by comparing the sample second moment of the elements in the putative target set, L , to those not in the set and also considering the product of the sample means of the two sets. The measurements of the target stimuli are equal and also the distractors measurements are similar when measurement correlations are perfect. Thus, on both target absent and present trials, F_L converges to a finite number, but the exponential prefactor in Eq. (7.12) diverges in the limit of $\rho_x \rightarrow 1$. Therefore, $d_{\text{NMT}}(\mathbf{x}) \rightarrow +\infty$ on target present trials and it goes to $-\infty$ on target absent trials. Hence, the observer performs perfectly in such a case. Further, the trend in performance for any

ρ_x when $\rho_s = 1$ can be explained in similar terms as seen in the case of a single target (Section 6.3.1.1).

In summary, we found that the impact of measurement correlations on the performance of an optimal Bayesian observer is significant with multiple targets as compared to the case of a single target (Chapter 6). This is attributed to the presence of more structured input and the interplay of measurement and stimulus correlations. Thus, presence of more structured displays enhances the joint effects of stimulus and measurement correlations.

However, the role of measurement correlations in response to the fixed input could also depend on the actual objective of the task rather than the true structure present in a visual scene. For instance, two different tasks can be performed with the same visual input but the impact of measurement correlations can vary extensively given the objective of the task. We consider an example of a discrimination task in the following section to explore the possible dependence of measurement correlations on the objective of a task.

7.4 Mean stimulus orientation discrimination task

As before, we consider that an observer is presented with N stimuli in a visual display. Stimulus orientations, denoted by $\mathbf{s} = (s_1, s_2, \dots, s_N)$, are relevant features of the task. We measure orientations relative to the vertical, which we denote by 0.

7.4. MEAN STIMULUS ORIENTATION DISCRIMINATION TASK

The task for the observer is to determine whether the mean orientation of the stimulus set is to the left (this is denoted by $C = -1$) or right ($C = 1$) of the vertical. The binary variable, C represents the two classes that need to be discriminated. Hence, this type of a task is known as a *discrimination task*.

Stimulus orientations are drawn from a multivariate normal distribution with mean vector, $\mathbf{0}_N$ and covariance matrix, $\Sigma_{\mathbf{s}}$, so that

$$\mathbf{s} \sim \mathcal{N}(\mathbf{0}_N, \Sigma_{\mathbf{s}}). \quad (7.13)$$

We have defined the matrix $\Sigma_{\mathbf{s}}$ in Eq. (2.3). The observer makes a decision based on the measurements of the presented stimuli, denoted by $\mathbf{x} = (x_1, x_2, \dots, x_N)$. Similar to the target detection tasks discussed in Sections 6.1 and 7.1, we assume the measurements to be unbiased and follow multivariate normal distribution with mean, \mathbf{s} and covariance matrix, $\Sigma_{\mathbf{x}}$ as

$$\mathbf{x}|\mathbf{s} \sim \mathcal{N}(\mathbf{s}, \Sigma_{\mathbf{x}}).$$

We observe that the framework of the task and thus the structure of visual inputs is similar to the target detection tasks analyzed in Chapter 6 and here in Section 7.1. But, we do not have any characterization of the target in this case. The observer is only interested in determining whether $C = 1$ or $C = -1$, instead of detecting any particular orientation of the stimuli.

7.4.1 Inference process

An optimal Bayesian observer needs to determine whether the mean orientation of the set of stimuli is oriented to the left or right of the vertical. We denote the mean orientation of the set of stimuli on a trial by \bar{s} , so that $\bar{s} = \sum_{i=1}^N s_i$. The optimal observer performs the computation based on the log-posterior ratio of the mean stimulus orientation being left or right, given the measurements. We denote the Bayesian decision variable on the task by $d_{\text{MD}}(\mathbf{x})$ and is given by

$$\begin{aligned} d_{\text{MD}}(\mathbf{x}) &= \log \frac{p(C = 1|\mathbf{x})}{p(C = -1|\mathbf{x})} = \log \frac{p(\bar{s} > 0|\mathbf{x})}{p(\bar{s} < 0|\mathbf{x})} \\ &= \log \frac{p(\mathbf{x}|\bar{s} > 0)}{p(\mathbf{x}|\bar{s} < 0)} + \log \frac{p(\bar{s} > 0)}{p(\bar{s} < 0)}. \end{aligned} \quad (7.14)$$

We denote the observer's MAP estimate of C by \hat{C} . When the decision variable, $d_{\text{MD}}(\mathbf{x}) > 0$, the observer infers $\hat{C} = 1$ and responds the mean stimulus orientation to the right of the vertical, that is, $\bar{s} > 0$. If $d_{\text{MD}}(\mathbf{x}) < 0$, the observer reports $\hat{C} = -1$, and $\bar{s} < 0$.

To compute the density function $p(\mathbf{x}|C)$ in Eq. (7.14), we marginalize the observer's information over the variable \mathbf{s} and further apply Bayes' rule to obtain

$$\begin{aligned} p(\mathbf{x}|\bar{s} > 0) &= \int p(\mathbf{x}|\mathbf{s})p(\mathbf{s}|\bar{s} > 0)d\mathbf{s} \\ &= \int p(\mathbf{x}|\mathbf{s})p(\bar{s} > 0|\mathbf{s})\frac{p(\mathbf{s})}{p(\bar{s} > 0)}d\mathbf{s} \\ &= \frac{1}{p(\bar{s} > 0)} \int_{\bar{s} > 0} p(\mathbf{x}|\mathbf{s})p(\mathbf{s})d\mathbf{s} \\ &= \frac{1}{p(\bar{s} > 0)} \int_{\bar{s} > 0} f(\mathbf{x}; \mathbf{s}, \Sigma_x) f(\mathbf{s}; \mathbf{0}_N, \Sigma_s) d\mathbf{s} \\ &= \frac{k_x}{p(\bar{s} > 0)} \int_{\bar{s} > 0} f\left(\mathbf{s}; \left(\mathbf{I} + \Sigma_x \Sigma_s^{-1}\right)^{-1} \mathbf{x}, \left(\Sigma_s^{-1} + \Sigma_x^{-1}\right)^{-1}\right) d\mathbf{s}. \end{aligned}$$

Appendix Eq. (B.3) is used above to simplify the product of two multivariate normal distributions and here $k_{\mathbf{x}} = \frac{1}{\sqrt{2\pi|\Sigma_{\mathbf{s}} + \Sigma_{\mathbf{x}}|}} \exp\left(-\frac{1}{2}\mathbf{x}^T(\Sigma_{\mathbf{s}} + \Sigma_{\mathbf{x}})^{-1}\mathbf{x}\right)$ is a normalization constant. Similarly, we compute $p(\mathbf{x}|\bar{s} < 0)$ and obtain the following expression for the Bayesian decision variable, $d_{\text{MD}}(\mathbf{x})$:

$$d_{\text{MD}}(\mathbf{x}) = \log \left[\frac{\int_{\bar{s} > 0} f\left(\mathbf{s}; \left(\mathbf{I} + \Sigma_{\mathbf{x}}\Sigma_{\mathbf{s}}^{-1}\right)^{-1}\mathbf{x}, \left(\Sigma_{\mathbf{s}}^{-1} + \Sigma_{\mathbf{x}}^{-1}\right)^{-1}\right) d\mathbf{s}}{\int_{\bar{s} < 0} f\left(\mathbf{s}; \left(\mathbf{I} + \Sigma_{\mathbf{x}}\Sigma_{\mathbf{s}}^{-1}\right)^{-1}\mathbf{x}, \left(\Sigma_{\mathbf{s}}^{-1} + \Sigma_{\mathbf{x}}^{-1}\right)^{-1}\right) d\mathbf{s}} \right]. \quad (7.15)$$

The above equation characterizes the decisions of the ideal observer on the task to discriminate whether the mean orientation of the presented stimuli is to the left or right of the vertical. Though the expression is not explicit, it depends on various parameters that govern the external structure of a visual scene and those that determine the structure in the observer's measurements. For instance, the total number of stimuli N , the variance and correlation of stimuli, σ_s^2 and ρ_s shape the structure of the input stimuli, while σ_x^2 and ρ_x determine the structure in the measurements.

We find that it is difficult to further simplify Eq. (7.15) analytically and therefore, we analyze the decisions of the ideal observer by performing numerical simulations. We present our simulation results below.

7.4.2 Results

We observe that the structure of input stimulus in the discrimination task is similar to the multiple target detection task discussed in Section 7.1. We are interested

7.4. MEAN STIMULUS ORIENTATION DISCRIMINATION TASK

in determining how the effects of measurement correlations change with the objective of a task. We thus analyze the impact of measurement correlations on the performance of an optimal observer in both tasks.

For comparison purposes, we consider the simulations of the mean stimulus orientation discrimination task with $N = 4$ stimuli, $\sigma_s = 15^\circ$ and $\sigma_x = 4^\circ$. Figure 7.2(A) shows the performance of an optimal Bayesian observer on the discrimination task as a function of ρ_x when $\rho_s = 0.5$.




Figure 7.2: Performance of an ideal Bayesian observer on the mean stimulus orientation discrimination task and multiple target detection task with $N = 4$ stimuli, $\sigma_s = 15^\circ$, $\sigma_x = 4^\circ$, and $\rho_s = 0.5$. (A) Mean stimulus orientation discrimination task. Performance of an optimal Bayesian observer as a function of measurement correlations, ρ_x on the mean stimulus orientation discrimination task. (B) Multiple target detection task. Similar to (A) in a multiple target detection task with $n = 2$.

We note that the performance gradually decreases as ρ_x increases. This is in contrast to the trend in performance on a multiple target detection task in Figure 7.2(B). The increasing measurement correlations improve the performance of

7.4. MEAN STIMULUS ORIENTATION DISCRIMINATION TASK

the ideal observer on the multiple target detection task and indeed perfect performance is possible when responses are perfectly correlated. However, measurement correlations have a negative impact on the performance when the task is to discriminate the mean stimulus orientation, instead of finding a target in similarly structured visual scenes.

Though the input is structured in a similar way in both cases, the observer needs to use different strategies to make decisions on the two tasks. In the discrimination task, the observer is required to integrate information from different sources; whereas the detection task requires extracting information about targets buried in a sea of distractors. Intuitively, a choice between two possibilities needs to be made based on the measurement of a set of stimuli. Each choice corresponds to a distribution of measurements. The difficulty of the task depends on how much the degree of overlap between these two distributions. The higher the overlap, the more difficult it is to tell which distribution a measurement belongs to, and the more difficult the decision. External structure, as well as structured noise in the measurements impacts the overlap between these distributions. Therefore, performance of an ideal observer depends not only on the strength, but also on the structure of measurement noise [6, 4]. Thus, the particulars of the task, the structure of the stimulus, as well as the strength and structure of the measurements jointly determine performance.

Hence, the role of measurement correlations can be subtle, and depends on the nature of the task. We thus have examined their impact on detecting a single (Chapter 6) and multiple targets among a group of distractors that have varying

7.4. MEAN STIMULUS ORIENTATION DISCRIMINATION TASK

degrees of pairwise correlations in their orientations.

Chapter 8

Discussion

Understanding how correlations between stimuli and measurements affect our decisions is important to understand how our visual perceptual system responds to structured input along with the structured measurements. Here we presented a thorough analysis of an experimental study designed to investigate how humans make decisions in response to stimuli that have varying degrees of structure on a visual search task. Further, we theoretically analyzed the joint effects of stimulus and measurement correlations on the performance of an ideal observer in a family of visual search tasks. Below, we discuss our findings and their limitations along with potential generalizations.

Stimulus correlations in visual search

Several recent studies found that humans perform near optimally on visual search tasks [94, 151, 98, 99]. That is, humans are capable of making best possible decisions while searching for a predefined target among non-relevant distractors. These studies only used two types of distractors - homogeneous (identical) and heterogeneous (independent). Moreover, the orientations of the homogeneous distractors were completely predictable (same across trials) except in [99]. Mazur et al. [99] were the first to distinguish homogeneity from predictability. They manipulated the statistics of the distractors by using trial-to-trial variability in homogeneous displays. The distractors were still identical to each other, but the orientation of the distractor was randomly chosen across trials. Regardless of introducing variability across trials in homogeneous displays and changing the degree of heterogeneity, their experiments were still focused on two extreme possibilities of distractors - with identical or independent random orientations.

Visual stimuli in natural scenes can be correlated to different extents with each other, therefore, it is important to understand how our visual decisions are affected by inputs that have intermediate correlations. In our work, we thus explored the intermediate regime of correlations between distractors. Using a target detection experiment, we examined whether humans take into account stimulus (distractors) correlations in visual search. We varied the amount of correlations between stimuli across different experimental sessions. Different correlations introduced varying amount of structure in visual displays - from none (uncorrelated

distractors), to mid-level (partially correlated distractors), to high-level (perfectly correlated distractors). We obtained mixed results based on the analysis of subjects' data and model selection. Due to individual differences, we did not obtain a clear presentation of how subjects inferred different correlation strengths. However, based on the assumptions of our best fitted models on different criteria, we found that subjects accounted for stimulus correlations in their responses, though, they were suboptimal in inferring the true correlation strength of distractor orientations in an experimental session.

Specifically, the favorable model based on Bayesian information criterion suggests that perhaps subjects were unable to distinguish the partially correlated conditions and they used a constant correlation strength to make decisions on those experimental sessions. While, subjects inferred a near-to-optimal correlation strength when distractors were perfectly correlated. This indicates that probably subjects use different inference processes when making decisions about homogeneous displays as compared to the case of partially structured displays. Perhaps humans perceive completely structured input differently while they may be unable to make a clear distinction between inputs having partial correlations and those having no structure. Further, we found that the distribution of encoding precision was dependent on the correlation strength used across experimental sessions. This dependence is difficult to explain, but suggests that possibly subjects encode stimuli with different level of precision when they are highly structured and differently when the statistical structure in a scene is relatively weak. A similar observation about different dependence of precision in homogeneous

and heterogeneous conditions was also made by Mazyar et al. [99]; however, they analyzed the set-size effects in visual search task.

The selected model based on Akaike information criterion suggests that subjects treated each experimental condition differently, but suboptimally. They incorrectly inferred the true correlation strength used to generate stimuli in an experimental session. Based on the parameter estimates of this model and experimentally obtained psychometric curves, we concluded that possibly subjects behave similarly in the conditions when distractors have intermediate strengths of correlations and they perform distinctly only in the case of identical distractors.

We note that these findings depend on the choice of models and the model comparison tools. It is always possible that there are better models and better explanations. For instance, we assumed that subjects correctly inferred the standard deviation of distractor orientations in the experiment. It might be possible that subjects used some other possible values of this standard deviation to make decisions. In that case, we need to use plausible assumptions on this parameter in our models and test whether subjects were able to infer this value correctly. It is also possible that subjects did not use any information about stimulus correlations or other parameters in the generative model of the task. Instead, they used alternative suboptimal strategies such as threshold criteria [94] to make their decisions. Besides these limitations about model assumptions, we are also limited by finite size of our data set. Our analysis and conclusions are based on the experimental data from a small number of participants. Therefore, individual subject differences are prominent in average model comparisons and lead to mixed results.

There are several other possibilities that exist and could perhaps provide a better explanation for the data. However, due to computational and time constraints, not every possibility can be tested. We tested a variety of models and these encompass a large range of possible assumptions about subjects' behavior. Thus, it is more likely that these models are sufficient to explain the decisions of human subjects on the experiment.

We examined a general question in our study. It is not only relevant to visual search studies, but to several other psychophysical studies in general. Different psychophysical experiments can be used to address the same question. For instance, the effects of structured environment on human decisions can be studied in a delayed estimation task [150]. Such a task could be simpler to perform as compared to the target detection task. In addition, a wide variety of stimuli are structured and correlated such as audio signals and odors. It is important to understand how brain integrates structured inputs to extract relevant information about the state of the world. Our work is a little step in this direction and we hope that our findings may have more general implications.

Interplay between stimulus and measurement correlations

Several theoretical and experimental studies have modeled decision processes by which the brain converts sensory measurements of a set of stimuli into a judgement about the world. But, many of them relied on stereotypical assumptions about the measurements being independent (across stimuli) and normally distributed [94, 98, 99, 151]. We extended our work here by focusing on the effects of violation of the assumption of independent measurements on performance in categorical, global perceptual judgements. It has been found that neural correlations can extend to long distances in visual cortex [39, 31], which suggests that the sensory measurements can be strongly correlated [124, 30]. But, the effects of measurement correlations cannot be studied in isolation; they need to be paired with the statistical structure of stimuli [99]. Therefore, in this work, we examined the joint influence of stimulus and measurement correlations on the performance of an ideal observer in a family of target detection tasks.

We found that the relation between stimulus and measurement correlations play a significant role in the decision-making strategy of an ideal Bayesian observer. Measurement correlations help in preserving the statistical structure of the stimuli on a multiple target detection task and hence enhance the performance; however, they have no effect on decisions in the case of a single target. This indicates that perhaps measurement correlations only play a role when sufficiently strong external structure is present in visual scenes. In the case of weak external

structure, visual displays are only partially structured and correlations between sensory measurements are inefficient in preserving those weak structures. We observed that the ideal observer always performs perfectly on both single and multiple target detection tasks when the distractors are homogeneous and measurements are perfectly correlated. Such an observation clearly reflects the significance of interaction between both correlations and the importance of measurement correlations in preserving the external statistical structure to improve the accuracy of decisions.

We also found that the influence of stimulus and measurement correlations depends on the relation between the external stimulus and measurement noise. In the regime of strong external noise as compared to the measurement noise, the trend in performance is easier to interpret in a single target detection task. While in the case of dominating measurement noise, we did not see a clear trend and it was relatively difficult to understand how both correlations interact to increase the performance at a particular value, but not at others.

Apart from the dependence on various parameters that control the statistical structure of visual displays, we found that influence of measurement correlations can also depend on the nature or objective of a task. Measurement correlations can have varying effects in response to same structured stimuli but with different task objectives. We observed that while these correlations enhance the performance of the ideal observer on a multiple target detection task, they negatively impact the decisions on a discrimination task. The statistical structure of visual scenes in both tasks were similar, but the observer was required to make two different

decisions - finding a predefined target among distractors and judging whether mean stimulus orientation of a set of stimuli orients to the left or right of vertical. This finding suggests that the role of measurement correlations can be subtle and the nature of the task along with the structure of stimulus jointly determine the influence of these correlations on performance of the observer in visual perceptual tasks.

We understand that our findings are limited to perceptual decisions in visual search tasks while the interaction between stimulus and measurement correlations can have a variety of other possible effects on our perceptual system. Therefore, it is important to understand how these correlations jointly interact, and govern our decisions. Further work needs to be done to identify the suitable parameter regimes where the effects of both these correlations are significant and can strongly impact our decision-making processes.

Appendix **A**

Notation table

Symbol	Description
$\mathcal{N}(\mu, \sigma^2)$	a one-dimensional normal distribution with mean μ and variance σ^2
$f(z; \mu, \sigma^2)$	a one-dimensional normal density function of variable, z with mean μ and variance, σ^2
$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	an N -dimensional multivariate normal distribution with N -dimensional mean vector $\boldsymbol{\mu}$ and $N \times N$ covariance matrix $\boldsymbol{\Sigma}$
$f(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	an N -dimensional normal density function of variable, \mathbf{z} with mean vector, $\boldsymbol{\mu}$ and covariance, $\boldsymbol{\Sigma}$; $\frac{1}{\sqrt{(2\pi)^N \boldsymbol{\Sigma} }} \exp\left(-\frac{1}{2}(\mathbf{z} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{z} - \boldsymbol{\mu})\right)$
$\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	a random variable \mathbf{z} having a normal distribution with probability density function, $f(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$
$\mathbf{z} Y \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	a random variable \mathbf{z} conditioned on Y having a normal distribution with probability density function, $f(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$
$\mathbf{0}_N$	an N -dimensional zero vector, $(0, 0, \dots, 0)$

Table A.1: *Continued on next page*

Table A.1: *Continued from previous page*

Symbol	Description
$\mathbf{1}_j$	an N -dimensional vector with j^{th} entry as 1 and rest zeros; $(0, 0, \dots, 0, 1, 0, \dots, 0)$
$\mathbf{z}_{\setminus j}$	$(z_1, z_2, \dots, z_{j-1}, z_{j+1}, \dots, z_N)$
$\Sigma_{\setminus j}$	an $N \times N$ matrix obtained by removing j^{th} row and column of matrix Σ
$ \mathbf{A} $	determinant of the matrix \mathbf{A}

Table A.1: **Mathematical Notations.** Description of the mathematical notations used in the text.

Appendix **B**

Some mathematical results

We present some known mathematical results and theorems here that are used in the main text to derive several results.

B.1 Product and integral of normal distributions

The product of m ($m \geq 2$) normal distributions over a single variable is a normal distribution. The integral of such a product distribution is analytically tractable. We list below few relevant results about the product and integral of normal distributions.

B.1.1 Univariate normal distributions

Consider two Gaussian probability distributions over a single random variable, z . Assume that the means of the two distributions are denoted by μ_1 and μ_2 ; and variances as σ_1^2 and σ_2^2 . The product of these two probability distributions is a normal distribution and the (normalized) probability density function of the resulting distribution can be computed as

$$\begin{aligned} f(z; \mu_1, \sigma_1^2) \cdot f(z; \mu_2, \sigma_2^2) &= f(\mu_1; \mu_2, \sigma_1^2 + \sigma_2^2) f\left(z; \frac{\frac{\mu_1}{\sigma_1^2} + \frac{\mu_2}{\sigma_2^2}}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}}, \frac{1}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}}\right) \\ &= k_c f(z; \mu_c, \sigma_c^2). \end{aligned} \quad (\text{B.1})$$

In general, the product of m such one-dimensional normal distributions over the same variable z is also a normal distribution. If the mean of the i^{th} distribution is denoted by μ_i and the variance by σ_i^2 , then the probability density function of the product of m such normal distributions is given by [93]

$$\prod_{i=1}^m f(z; \mu_i, \sigma_i^2) = c_p f\left(z; \frac{\sum_{i=1}^m \mu_i}{\sum_{i=1}^m \frac{1}{\sigma_i^2}}, \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}}\right), \quad (\text{B.2})$$

where $c_p = \frac{1}{\left(\prod_{i=1}^m \sigma_i\right) \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^2}}} \exp\left(-\frac{1}{2} \left(\sum_{i=1}^m \frac{\mu_i^2}{\sigma_i^2} - \frac{\left(\sum_{i=1}^m \frac{\mu_i}{\sigma_i}\right)^2}{\sum_{i=1}^m \frac{1}{\sigma_i^2}}\right)\right)$ is a normalization constant.

We now describe the results for the product and integral of the resulting product in the case of multivariate normal distributions.

B.1.2 Multivariate normal distributions

Consider two N -dimensional multivariate normal distributions in a single random variable, \mathbf{z} with mean vectors $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$; and covariance matrices, $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_2$. Given that the covariance matrices are nonsingular, the product of these two N -dimensional multivariate normal distributions in the random variable \mathbf{z} is another (unnormalized) N -dimensional multivariate normal distribution [2, 115] given by

$$f(\mathbf{z}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \cdot f(\mathbf{z}; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) = k_p f(\mathbf{z}; \boldsymbol{\mu}_p, \boldsymbol{\Sigma}_p), \quad (\text{B.3})$$

with

$$\begin{aligned} \boldsymbol{\Sigma}_p &= (\boldsymbol{\Sigma}_1^{-1} + \boldsymbol{\Sigma}_2^{-1})^{-1}, \quad \boldsymbol{\mu}_p = \boldsymbol{\Sigma}_p(\boldsymbol{\Sigma}_1^{-1}\boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_2^{-1}\boldsymbol{\mu}_2), \quad \text{and} \\ k_p &= |2\pi\boldsymbol{\Sigma}_1\boldsymbol{\Sigma}_2\boldsymbol{\Sigma}_p^{-1}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \boldsymbol{\Sigma}_1^{-1} \boldsymbol{\Sigma}_p \boldsymbol{\Sigma}_2^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)\right) \\ &= |2\pi(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)\right). \end{aligned}$$

In the case of one-dimensional normal distributions, Eq. (B.3) reduces to Eq. (B.1) with k_p reducing to $k_c = f(\mu_1; \mu_2, \sigma_1^2 + \sigma_2^2)$.

Further, since $\int f(\mathbf{z}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) d\mathbf{z} = 1$, the above results for the product of normal distributions implies that

$$\begin{aligned} \text{N-dimensional:} \quad & \int_{\mathbb{R}^N} f(\mathbf{z}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \cdot f(\mathbf{z}; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) d\mathbf{z} = \int_{\mathbb{R}^N} k_p f(\mathbf{z}; \boldsymbol{\mu}_p, \boldsymbol{\Sigma}_p) d\mathbf{z} = k_p \\ \text{one-dimensional:} \quad & \int_{\mathbb{R}} f(z; \mu_1, \sigma_1^2) \cdot f(z; \mu_2, \sigma_2^2) dz = \int_{\mathbb{R}} k_c f(z; \mu_c, \sigma_c^2) dz = k_c. \end{aligned} \quad (\text{B.4})$$

B.2 Determinant and inverse of a rank-1 matrix

The determinant of a rank-1 matrix can be computed using the following matrix determinant lemma [64, 25].

Lemma 1 (Matrix Determinant Lemma). *Suppose \mathbf{A} is an invertible square matrix, and \mathbf{u}, \mathbf{v} are column vectors. Then*

$$|\mathbf{A} + \mathbf{u}\mathbf{v}^T| = (1 + \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}) |\mathbf{A}|.$$

Here $\mathbf{u}\mathbf{v}^T$ is the outer product of two vectors u , and v .

Further, the inverse of a rank-1 matrix can be obtained using the Sherman-Morrison formula [38, 58, 138, 9, 59, 102, 118]. This is a special case of the generalized Woodbury formula given in Theorem 2.

Theorem 1 (Sherman-Morrison formula). *Suppose \mathbf{A} is an invertible square matrix, and \mathbf{u}, \mathbf{v} are vectors. Assume that $1 + \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u} \neq 0$. Then*

$$(\mathbf{A} + \mathbf{u}\mathbf{v}^T)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{A}^{-1}}{1 + \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}}.$$

Here $\mathbf{u}\mathbf{v}^T$ is the outer product of two vectors u , and v .

B.3 Determinant and inverse of a rank- k matrix

We now present the relevant general theorems for computing the determinant and inverse of a rank- k matrix.

B.3. DETERMINANT AND INVERSE OF A RANK-K MATRIX

Lemma 2 (Generalized Matrix Determinant Lemma). *Suppose \mathbf{A} is an invertible $n \times n$ matrix, \mathbf{U}, \mathbf{V} are $n \times m$ matrices, and \mathbf{W} is an invertible $m \times m$ matrix. Then*

$$|\mathbf{A} + \mathbf{U}\mathbf{W}\mathbf{V}^T| = |\mathbf{W}^{-1} + \mathbf{V}^T\mathbf{A}^{-1}\mathbf{U}||\mathbf{W}||\mathbf{A}|.$$

The inverse of a rank- k matrix can be computed using the following generalized Woodbury matrix identity [163, 59, 66].

Theorem 2 (Sherman-Morrison Woodbury formula or Woodbury matrix identity). *Suppose \mathbf{A} is an invertible $n \times n$ matrix, \mathbf{U}, \mathbf{W} , and \mathbf{V} are $n \times m$, $m \times m$, and $m \times n$ matrices. Then*

$$(\mathbf{A} + \mathbf{U}\mathbf{W}\mathbf{V})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U}(\mathbf{W}^{-1} + \mathbf{V}\mathbf{A}^{-1}\mathbf{U})^{-1}\mathbf{V}\mathbf{A}^{-1}.$$

Gabor Filter

A Gabor filter is a linear filter having frequency and orientation similar to that of human visual system, and is obtained by multiplication of a sinusoidal wave with Gaussian kernel function [45, 104]. It can be used as a band-pass filter for unidimensional signals (e.g. speech). The filter has a real and imaginary component given by

$$g(x, y; \lambda, \theta, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cos\left(2\pi\frac{x'}{\lambda} + \psi\right),$$

and

$$g(x, y; \lambda, \theta, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \sin\left(2\pi\frac{x'}{\lambda} + \psi\right).$$

Here $x' = x \cos \theta + y \sin \theta$ and $y' = -x \sin \theta + y \cos \theta$. In the above equations, λ represents the wavelength of the sinusoidal wave, θ is the angle of the normal to the parallel stripes of a Gabor function, ψ is the phase offset, σ is the standard deviation of the Gaussian envelope, and γ denotes the spatial aspect ratio and describes the ellipticity of the support of the Gabor function.

Experimental studies have shown that simple cells in human visual system can be modeled by Gabor functions [33, 70]. Thus, Gabor filters are extensively used as stimuli in psychophysical studies. An example of a Gabor patch is shown in Figure 2.2(A).

Bibliography

- [1] L.F. Abbott and P. Dayan. The effect of correlated variability on the accuracy of a population code. *Neural Computation*, 11(1):91–101, 1999.
- [2] P. Ahrendt. The multivariate Gaussian probability distribution. Technical report, Technical University of Denmark, 2005. <http://www2.imm.dtu.dk/pubdb/p.php?3312>.
- [3] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.
- [4] B.B. Averbeck. *Noise correlations and information encoding and decoding*. Springer, 2009.
- [5] B.B. Averbeck and D. Lee. Effects of noise correlations on information encoding and decoding. *Journal of Neurophysiology*, 95(6):3633–3644, 2006.
- [6] B.B. Averbeck, P.E. Latham, and A. Pouget. Neural correlations, population coding and computation. *Nature Neuroscience Reviews*, 7(5):358–366, 2006.
- [7] S. Baldassi and D.C. Burr. Feature-based integration of orientation signals in visual search. *Vision Research*, 40(10):1293–1300, 2000.
- [8] S. Baldassi and P. Verghese. Comparing integration rules in visual search. *Journal of Vision*, 2(8):559–70, 2002.
- [9] M.S. Bartlett. An inverse matrix adjustment arising in discriminant analysis. *The Annals of Mathematical Statistics*, 22(1):107–111, 1951.

- [10] P.W. Battaglia, R.A. Jacobs, and R.N. Aslin. Bayesian integration of visual and auditory signals for spatial localization. *Journal of the Optical Society of America A*, 20(7):1391–1397, 2003.
- [11] W.F. Bauer. The Monte Carlo method. *Journal of the Society for Industrial & Applied Mathematics*, 6(4):438–451, 1958.
- [12] T. Bayes and R. Price. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions (1683-1775)*, pages 370–418, 1763.
- [13] J. Beck, J. Kanitscheider, X. Pitkow, P. Latham, and A. Pouget. The perils of inferring information from correlations. In *Cosyne Abstracts 2013, Salt Lake City USA*, 2013.
- [14] J.M. Beck, W.J. Ma, X. Pitkow, P.E. Latham, and A. Pouget. Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, 74(1):30–39, 2012.
- [15] B.A. Berg. *Markov chain Monte Carlo simulations and their statistical analysis: with web-based Fortran code*. World Scientific Publishers, Singapore, 2004.
- [16] M. Bhardwaj, R. van den Berg, W.J. Ma, and Krešimir Josić. Do humans account for stimulus correlations in visual perception? In *Cosyne Abstracts 2013, Salt Lake City, USA*, 2013.
- [17] M. Bhardwaj, R. van den Berg, W.J. Ma, and Krešimir Josić. Visual search in structured environments. *In preparation*, 2013.
- [18] H.S. Bhat and N. Kumar. On the derivation of the Bayesian information criterion. School of Natural Sciences, University of California. <http://nscs00.ucmerced.edu/~nkumar4/BhatKumarBIC.pdf>, 2010.
- [19] C.M. Bishop and N.M. Nasrabadi. *Pattern recognition and machine learning*, volume 1. Springer, New York, 2006.
- [20] G.E.P. Box. Evolutionary operation: A method for increasing industrial productivity. *Applied Statistics*, 6(2):81–101, 1957.
- [21] H. Bozdogan. Model selection and Akaike’s information criterion (AIC): The general theory and its analytical extensions. *Psychometrika*, 52(3):345–370, 1987.
- [22] H. Bozdogan. Akaike’s information criterion and recent developments in information complexity. *Journal of Mathematical Psychology*, 44(1):62–91, 2000.

- [23] T.F. Brady and J.B. Tenenbaum. Encoding higher-order structure in visual working memory: A probabilistic model. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, pages 411–416, 2010.
- [24] D.H. Brainard. The psychophysics toolbox. *Spatial Vision*, 10(4):433–436, 1997.
- [25] M. Brookes. The matrix reference manual. <http://www.ee.imperial.ac.uk/hp/staff/dmb/matrix/intro.html>, 2011.
- [26] V. Bruce. *Visual perception: Physiology, psychology, and ecology*. Psychology Press, 2003.
- [27] K.P. Burnham and D.R. Anderson. Multimodel inference understanding AIC and BIC in model selection. *Sociological Methods & Research*, 33(2):261–304, 2004.
- [28] S. Carroll, M. Bhardwaj, W.J. Ma, and K. Josić. Visual decisions in presence of external and internal correlations. *in preparation*, 2013.
- [29] L.K.H. Chan and W.G. Hayward. Visual search. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2013.
- [30] Y. Chen, W.S. Geisler, and E. Seidemann. Optimal decoding of correlated neural population responses in the primate visual cortex. *Nature Neuroscience*, 9(11):1412–1420, 2006.
- [31] M.R. Cohen and A. Kohn. Measuring and interpreting neuronal correlations. *Nature Neuroscience*, 14(7):811–819, 2011.
- [32] C.A. Colin and F.A.G. Windmeijer. An R-squared measure of goodness of fit for some common nonlinear regression models. *Journal of Econometrics*, 77(2):329–342, 1997.
- [33] J.G. Daugman. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Optical Society of America, Journal, A: Optics and Image Science*, 2(7):1160–1169, 1985.
- [34] P. Dayan and L. Abbott. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT Press, 2001.
- [35] P. Dayan, G.E. Hinton, R.M. Neal, and R.S. Zemel. The Helmholtz machine. *Neural Computation*, 7(5):889–904, 1995.

- [36] J. de la Rocha, B. Doiron, E. Shea-Brown, K. Josić, and A. Reyes. Correlation between neural spike trains increases with firing rate. *Nature*, 448(7155):802–806, 2007.
- [37] J. Duncan and G.W. Humphreys. Visual search and stimulus similarity. *Psychological Review*, 96(3):433, 1989.
- [38] W.J. Duncan. Some devices for the solution of large sets of simultaneous linear equations (with an appendix on the reciprocation of partitioned matrices). *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, Seventh Series*, 35:660–670, 1944.
- [39] A.S. Ecker, P. Berens, G.A. Keliris, M. Bethge, N.K. Logothetis, and A.S. Tolias. Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965):584–587, 2010.
- [40] A.S. Ecker, P. Berens, A.S. Tolias, and M. Bethge. The effect of noise correlations in populations of diversely tuned neurons. *The Journal of Neuroscience*, 31(40):14272–14283, 2011.
- [41] M.P. Eckstein. Visual search: A retrospective. *Journal of Vision*, 11(5):1–36, 2011.
- [42] M.P. Eckstein, J.P. Thomas, J. Palmer, and S.S. Shimozaki. A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays. *Perception & Psychophysics*, 62(3):425–451, 2000.
- [43] M.O. Ernst and M.S. Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, 2002.
- [44] W.K. Estes and H.A. Taylor. A detection method and probabilistic models for assessing information processing from brief visual displays. *Proceedings of the National Academy of Sciences of the United States of America*, 52(2):446, 1964.
- [45] I. Fogel and D. Sagi. Gabor filters as texture discriminator. *Biological Cybernetics*, 61(2):103–113, 1989.
- [46] D. Fougnie, J.W. Suchow, and G.A. Alvarez. Variability in the quality of visual working memory. *Nature Communications*, 3:1229, 2012.
- [47] F. Gabbiani and S.J. Cox. *Mathematics for neuroscientists*. Academic Press, 2010.

- [48] E. Ganmor, R. Segev, and E. Schneidman. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proceedings of the National Academy of Sciences*, 108(23):9679–9684, 2011.
- [49] T.J. Gawne and B.J. Richmond. How independent are the messages carried by adjacent inferior temporal cortical neurons? *The Journal of Neuroscience*, 13(7):2758–2771, 1993.
- [50] A. Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin. *Bayesian data analysis*. Chapman & Hall, CRC press, 2003.
- [51] G.A. Gescheider. *Psychophysics: the fundamentals*. Psychology Press, 1997.
- [52] D.E. Goldberg and J.H. Holland. Genetic algorithms and machine learning. *Machine Learning*, 3(2):95–99, 1988.
- [53] D. Goldreich. A Bayesian perceptual model replicates the cutaneous rabbit and other tactile spatiotemporal illusions. *PLoS One*, 2(3):e333, 2007.
- [54] D. Goldreich and M.A. Peterson. A Bayesian observer replicates convexity context effects in figure-ground perception. *Seeing and Perceiving*, 25(3-4):3–4, 2012.
- [55] N. Graham, P. Kramer, and D. Yager. Signal-detection models for multidimensional stimuli: Probability distributions and combination rules. *Journal of Mathematical Psychology*, 31(4):366–409, 1987.
- [56] D.M. Green and J.A. Swets. *Signal detection theory and psychophysics*, volume 1. John Wiley & Sons, New York, 1966.
- [57] D.A. Gutnisky and V. Dragoi. Adaptive coding of visual information in neural populations. *Nature*, 452(7184):220–224, 2008.
- [58] L. Guttman. Enlargement methods for computing the inverse matrix. *The Annals of Mathematical Statistics*, 17(3):336–343, 1946.
- [59] W.H. Hager. Updating the inverse of a matrix. *SIAM Review*, 31(2):221–239, 1989.
- [60] J.H. Halton. A retrospective and prospective survey of the Monte Carlo method. *SIAM Review*, 12(1):1–63, 1970.
- [61] W. Hamilton. *Lectures on metaphysics and logic*, volume 1. Gould and Lincoln, 1859.

- [62] J.M. Hammersley and D.C. Handscomb. *Monte Carlo methods*. Chapman and Hall, 1964.
- [63] B.J. Hansen, M.I. Chelaru, and V. Dragoi. Correlated variability in laminar cortical circuits. *Neuron*, 76:590–602, 2012.
- [64] D.A. Harville. Matrix algebra from a statistician’s perspective. *Technometrics*, 40(2):164–164, 1998.
- [65] G. Hatfield. Perception as unconscious inference. *Perception and the physical world: Psychological and philosophical issues in perception*, pages 113–143, 2002.
- [66] N.J. Higham. *Accuracy and stability of numerical algorithms*. SIAM, 1996.
- [67] R. Hooke and T.A. Jeeves. Direct search solution of numerical and statistical problems. *Journal of the ACM (JACM)*, 8(2):212–229, 1961.
- [68] R. A. Jacobs. Optimal integration of texture and motion cues to depth. *Vision Research*, 39(21):3621–3629, 1999.
- [69] E.T. Jaynes. *How does the brain do plausible reasoning?* Springer, 1988.
- [70] J.P. Jones and L.A. Palmer. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.
- [71] K. Josić, E. Shea-Brown, B. Doiron, and J. de la Rocha. Stimulus-dependent correlations and population codes. *Neural Computation*, 21(10):2774–2804, 2009.
- [72] M.H. Kalos and P.A. Whitlock. *Monte Carlo methods*. John Wiley & Sons, 2008.
- [73] S.M. Kay. *Fundamentals of statistical signal processing, Detection Theory, volume 2*. Prentice Hall PTR, New Jersey, USA, 1998.
- [74] D. Kersten, P. Mamassian, and A. Yuille. Object perception as Bayesian inference. *Annual Review of Psychology*, 55:271–304, 2004.
- [75] S. Keshvari, R. van den Berg, and W.J. Ma. Probabilistic computation in human perception under variability in encoding precision. *PLoS ONE*, 7(6):e40216, 2012.
- [76] S. Keshvari, R. van den Berg, and W.J. Ma. No evidence for an item limit in change detection. *PLoS Computational Biology*, 9(2):e1002927, 2013.

- [77] S.A. Klein. Measuring, estimating, and understanding the psychometric function: A commentary. *Perception & Psychophysics*, 63(8):1421–1455, 2001.
- [78] D.C. Knill. Reaching for visual cues to depth: The brain combines depth cues differently for motor control and perception. *Journal of Vision*, 5(2):103–115, 2005.
- [79] D.C. Knill and A. Pouget. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004.
- [80] D.C. Knill and W. Richards. *Perception as Bayesian inference*. Cambridge University Press, 1996.
- [81] A. Kohn and M.A. Smith. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *The Journal of Neuroscience*, 25(14):3661–3673, 2005.
- [82] K. Körding. Decision theory: what should the nervous system do? *Science*, 318(5850):606–610, 2007.
- [83] K.P. Körding, S. Ku, and D.M. Wolpert. Bayesian integration in force estimation. *Journal of Neurophysiology*, 92(5):3161–3165, 2004.
- [84] K.P. Körding and D.M. Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247, 2004.
- [85] S. Kullback and R.A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [86] P.E. Latham and S. Nirenberg. Synergy, redundancy, and independence in population codes, revisited. *The Journal of Neuroscience*, 25(21):5195–5206, 2005.
- [87] P.E. Latham and Y. Roudi. Role of correlations in population coding. *arXiv preprint:1109.6524*, 2011.
- [88] P.M. Lee. *Bayesian statistics: an introduction*. John Wiley & Sons, 2012.
- [89] R. Lewis and V. Torczon. Pattern search algorithms for bound constrained minimization. *SIAM Journal on Optimization*, 9(4):1082–1099, 1999.
- [90] W.J. Ma. Bayesian approach to perception. *SAGE Encyclopedia of Perception*, pages 201–205, 2009.

- [91] W.J. Ma. Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Research*, 50(22):2308–2319, 2010.
- [92] W.J. Ma. Theoretical neuroscience-learning, perception and cognition. Course Lecture notes, Spring 2011.
- [93] W.J. Ma, K. Körding, and D. Goldreich. Bayesian modeling of perception. unpublished, 2013.
- [94] W.J. Ma, V. Navalpakkam, J. M. Beck, R. van den Berg, and A. Pouget. Behavior and neural basis of near-optimal visual search. *Nature Neuroscience*, 14(6):783–790, 2011.
- [95] D.J.C. MacKay. Bayesian interpolation. *Neural Computation*, 4(3):415–447, 1992.
- [96] D.J.C. MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.
- [97] N.A. Macmillan and D.C. Creelman. *Detection theory: A user's guide*. Psychology Press, 2004.
- [98] H. Mazyar, R. van den Berg, and W.J. Ma. Does precision decrease with set size? *Journal of Vision*, 12(6):1–16, 2012.
- [99] H. Mazyar, R. van den Berg, R.L. Seilheimer, and W.J. Ma. Independence is elusive: set size effects on encoding precision in visual search. *Journal of Vision*, 13(5):1–14, 2013.
- [100] J.L. McClelland. Connectionist models and Bayesian inference. *Rational Models of Cognition*, pages 21–53, 1998.
- [101] D. McNicol. *A primer of signal detection theory*. Psychology Press, 2005.
- [102] C.D. Meyer. *Matrix analysis and Applied Linear Algebra Book and Solutions Manual*, volume 2. SIAM, 2000.
- [103] M. Mitchell and S. Forrest. Genetic algorithms and artificial life. *Artificial Life*, 1(3):267–289, 1994.
- [104] J.R. Movellan. Tutorial on gabor filters. <http://www.ai.polymt1.ca/cohen/ELE6812/Gabor.pdf>, 2008.
- [105] A.L. Nagy and G. Thomas. Distractor heterogeneity, attention, and color in visual search. *Vision Research*, 43(14):1541–1552, 2003.

- [106] V. Navalpakkam and L. Itti. An integrated model of top-down and bottom-up attention for optimizing detection speed. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 2049–2056. IEEE, 2006.
- [107] R.M. Neal. Probabilistic inference using Markov chain Monte Carlo methods. Technical report, Department of Computer Science, University of Toronto, Toronto, Ontario, Canada, 1993. <http://omega.albany.edu:8008/neal.pdf>.
- [108] L.W. Nolte and D. Jaarsma. More on the detection of one of m orthogonal signals. *The Journal of the Acoustical Society of America*, 41:497, 1967.
- [109] I.F. Ohiorhenuan, F. Mechler, K.P. Purpura, A.M. Schmid, Q. Hu, and J.D. Victor. Sparse coding and high-order correlations in fine-scale cortical network. *Nature*, 466:617–621, 2010.
- [110] J. Palmer. Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology: Human Perception and Performance*, 16(2):332–350, 1990.
- [111] J. Palmer. Set-size effects in visual search: The effect of attention is independent of the stimulus for simple tasks. *Vision Research*, 34(13):1703–1721, 1994.
- [112] J. Palmer, P. Verghese, and M. Pavel. The psychophysics of visual search. *Vision Research*, 40(10):1227–1268, 2000.
- [113] D.G. Pelli. The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4):437–442, 1997.
- [114] D.H. Perkel, G.L. Gerstein, and G.P. Moore. Neuronal spike trains and stochastic point processes: II. Simultaneous spike trains. *Biophysical Journal*, 7(4):419–440, 1967.
- [115] K.B. Petersen and M.S. Pedersen. The matrix cookbook. Technical University of Denmark. <http://www2.imm.dtu.dk/pubdb/p.php?3274>, 2008.
- [116] W.W. Peterson, T.G. Birdsall, and W. Fox. The theory of signal detectability. *Information Theory, IRE Professional Group on*, 4(4):171–212, 1954.
- [117] A. Pouget, J.M. Beck, W.J. Ma, and P.E. Latham. Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9):1170–1178, 2013.

- [118] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge University Press, 2007.
- [119] R.P.N. Rao and D.H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, 1999.
- [120] R.P.N. Rao, B.A. Olshausen, and M.S. Lewicki. *Probabilistic models of the brain: Perception and neural function*. MIT Press, 2002.
- [121] C.P. Robert and G. Casella. *Monte Carlo statistical methods*, volume 319. Springer-Verlag, 2004.
- [122] R. Romo, A. Hernandez, A. Zainos, and E. Salinas. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron*, 38(4):649–657, 2003.
- [123] R. Rosenbaum and K. Josić. Mechanisms that modulate the transfer of spiking correlations. *Neural Computation*, 23(5):1261–1305, 2011.
- [124] R. Rosenbaum, J. Trousdale, and Krešimir Josić. Pooling and correlated neural activity. *Frontiers in Computational Neuroscience*, 4(9):1–14, 2010.
- [125] R. Rosenholtz. Visual search for orientation among heterogeneous distractors: experimental results and implications for signal-detection theory models of search. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4):985–999, 2001.
- [126] R.Y. Rubinstein and D.P. Kroese. *Simulation and the Monte Carlo method*. Wiley Series in Probability and Mathematical Statistics, New York, 1981.
- [127] E. Salinas and T.J. Sejnowski. Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(8):539–550, 2001.
- [128] L.L. Scharf. *Statistical signal processing*, volume 98. Addison-Wesley Reading, 1991.
- [129] E. Schneidman, M. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

- [130] E. Schneidman, W. Bialek, and M.J. Berry. Synergy, redundancy, and independence in population codes. *The Journal of Neuroscience*, 23(37):11539–11553, 2003.
- [131] G.E. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.
- [132] P. Seriès, P.E. Latham, and A. Pouget. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nature Neuroscience*, 7(10):1129–1135, 2004.
- [133] M. Shamir and H. Sompolinsky. Correlation codes in neuronal populations. *Advances in Neural Information Processing Systems*, 1:277–284, 2002.
- [134] L. Shams, W.J. Ma, and U. Beierholm. Sound-induced flash illusion as an optimal percept. *Neuroreport*, 16(17):1923–1927, 2005.
- [135] M.L. Shaw. Identifying attentional and decision-making components in information processing. *Attention and Performance VIII*, 8:277–295, 1980.
- [136] M.L. Shaw. Division of attention among spatial locations: A fundamental difference between detection of letters and detection of luminance increments. *Attention and Performance X*, pages 109–121, 1984.
- [137] S. Shen, R. van den Berg, and W.J. Ma. When is sensory precision variable? In *Cosyne Abstracts 2013, Salt Lake City, USA*, 2013.
- [138] J. Sherman and W.J. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950.
- [139] H. Shin, R. van den Berg, and W.J. Ma. Independent pools of visual short-term memory resource for different features. In *Cosyne Abstracts 2013, Salt Lake City, USA*, 2013.
- [140] J.D. Smith, J.S. Redford, L.C. Gent, and D.A. Washburn. Visual search and the collapse of categorization. *Journal of Experimental Psychology: General*, 134(4):443–460, 2005.
- [141] M.A. Smith. *Alhacen’s theory of visual perception: A critical edition, with english translation and commentary, of the first three books of Alhacen’s De Aspectibus, the Medieval Latin version of Ibn Al-Haytham’s Kitāb Al-Manazir*, volume 1. American Philosophical Society, 2001.

- [142] H. Sompolinsky, H. Yoon, K. Kang, and M. Shamir. Population coding in neuronal systems with correlated noise. *Physical Review E*, 64(5):051904, 2001.
- [143] D. L. Swets and B. Punch. Genetic algorithms for object localization in a complex scene. In *IEEE International Conference on Image Processing*, pages 595–598. IEEE, 1995.
- [144] H. Tassinari, T.E. Hudson, and M.S. Landy. Combining priors and noisy visual cues in a rapid pointing task. *The Journal of Neuroscience*, 26(40):10154–10163, 2006.
- [145] G. Tkačik, J.S. Prentice, V. Balasubramanian, and E. Schneidman. Optimal population coding by noisy spiking neurons. *Proceedings of the National Academy of Sciences*, 107(32):14419–14424, 2010.
- [146] V. Torczon. On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25, 1997.
- [147] A. Treisman. Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance*, 8(2):194–214, 1982.
- [148] A.M. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136, 1980.
- [149] R. van den Berg, E. Awh, and W.J. Ma. Factorial comparison of working memory models. *Psychological Review*, in press, 2014.
- [150] R. van den Berg, H. Shin, W.C. Chou, R. George, and W.J. Ma. Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, 109(22):8780–8785, 2012.
- [151] R. van den Berg, M. Vogel, K. Josić, and W.J. Ma. Optimal inference of sameness. *Proceedings of the National Academy of Sciences*, 109(8):3178–3183, 2012.
- [152] P. Verghese. Visual search and attention: A signal detection theory approach. *Neuron*, 31(4):523–535, 2001.
- [153] B. T. Vincent, R. J. Baddeley, T. Troscianko, and I.D. Gilchrist. Optimal feature integration in visual search. *Journal of Vision*, 9(5):1–11, 2009.

- [154] H. von Helmholtz. Handbuch der Physiologischen Optik, vol. II. In G. Karsten (Ed.) *Allgemeine Encyklopedie der Physik*, 9:37–51, 1867.
- [155] L. Wasserman. Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, 44(1):92–107, 2000.
- [156] A.B. Watson and D.G. Pelli. Quest: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, 33(2):113–120, 1983.
- [157] F.A. Wichmann and N.J. Hill. The psychometric function: I. fitting, sampling, and goodness of fit. *Perception & Psychophysics*, 63(8):1293–1313, 2001.
- [158] T.D. Wickens. *Elementary signal detection theory*. Oxford university press, 2002.
- [159] P. Wilken and W.J. Ma. A detection theory account of change detection. *Journal of Vision*, 4(12):1120–1135, 2004.
- [160] R.L. Winkler. *An introduction to Bayesian inference and decision*. Holt, Rinehart and Winston New York, 1972.
- [161] J.M. Wolfe, M.L.H. Võ and K.K. Evans, and M.R. Greene. Visual search in scenes involves selective and nonselective pathways. *Trends in Cognitive Sciences*, 15(2):77–84, 2011.
- [162] J.M. Wolfe, S.R. Friedman-Hill, M.I. Stewart, and K.M. O’Connell. The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 18(1):34–49, 1992.
- [163] M.A. Woodbury. Inverting modified matrices. Technical report, Statistical Research Group, Princeton University, 1950.
- [164] W. Zhang and S.J. Luck. Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192):233–235, 2008.