POLITECNICO DI TORINO

Master's Degree in Physics of Complex Systems

Master's Degree Thesis

Stochastic Resonance in Human Sensory Processing

Supervisors

Prof. Peter NERI

Prof. Alessandro PELIZZOLA

Candidate

CARLO PARIS

April 2022

Abstract

In this work we give a brief overview of some fundamental concepts of Signal Detection Theory, an established framework in psychophysics. After that, we introduce the phenomenon of stochastic resonance, a non-linear effect whereby signal detectability is enhanced by addition of small amounts of uncorrelated noise. While introducing them, we concentrate mostly on concepts relevant for our research. Whether stochastic resonance is a legitimate phenomenon within the human brain or merely a secondary phenomenon is still largely a matter of dispute. Subsequently, we outline our motivations for studying stochastic resonance effect in the human brain and propose which controls should be applied for obtaining conclusive evidence to settle the aforementioned controversy. Finally, we present preliminary results of our pilot experiments and conclude by proposing potential future directions of research.

Summary

TITLE: Stochastic Resonance in Human Sensory Processing
CANDIDATE: Carlo Paris
SUPERVISORS: prof. Peter Neri, prof. Alessandro Pelizzola

This master's thesis is motivated by two objectives: a critical discussion of current evidence supporting stochastic resonance as a genuine phenomenon in human sensory processing; the active contribution of novel experimental evidence specifically designed to address issues left unresolved by existing data.
Stochastic resonance (SR) is a nonlinear phenomenon whereby a sub-threshold signal is rendered supra-threshold by the addition of small amounts of uncorrelated noise. This phenomenon is generally attributed to the contribution of noise components that fall within the energy band occupied by the signal: these components may resonate with the signal and push detector activation above threshold.

SR can be studied quantitatively in humans using psychophysical techniques. The primary goal of psychophysics is to quantify and model subjective experience in relation to external stimuli. The theoretical foundations of sensory psychophysics are formulated within an influential framework known as Signal Detection Theory (SDT), a statistical theory that models how agents respond to input stimuli when solving well-defined tasks. Minimal requirements for a psychophysical experiment are (1) a task for the human subject to perform with a well-defined goal (e.g. detect target X), (2) at least 2 possible states of the world (e.g. target X is present versus target X is absent), (3) one or multiple presentation intervals during which the subject is presented with sensory information regarding the state of the world and, finally, (4) the subject's behavioural output communicating their decision about the state of the world. One cycle containing these 4 ingredients is called a 'trial.'

A typical psychophysical experiment involves two visual stimuli: a known signal superimposed onto a noise background (corresponding to state-of-the-world *sn*); the noise background alone(state *n*). In the Yes-No protocol (YN), the human subject only sees one stimulus or the other and is asked to report the inferred state of the world (*sn* versus *n*). This task can be formulated to the subject with the question: was the signal present or absent? In the 2AFC protocol, the subject is

always presented with both stimuli: one in the *sn* state, and one in the *n* state; they are then asked to decide which stimulus was in state *sn* and which stimulus was in state *n*. This task can be formulated to the subject with the question: which stimulus contained the signal? The goal is to identify the correct state of the world on as many trials as possible.

From the viewpoint of SDT, the perceptual system constructs an internal representation of the probability distributions associated with the possible states of the world. These, combined with a decision criterion, support a behavioural decision. We can formalize this viewpoint using the language of Bayesian statistics. In general, we call prior the probability that the j^{th} state of the world is the actual state, and denote it $p(h_j)$. Secondly, the posterior is the probability that the state of the world j is the real one, given that some evidence e has been observed, mathematically expressed as $p(h_j|e)$. Here, e is the information provided to the subject regarding the true state of the world. Thirdly, the likelihood is written as $p(e|h_i)$, and is the probability that evidence e supports state of the world h_i . We can use likelihood to compute the likelihood ratio:

$$l_{ij}(e) = \frac{p(e|h_i)}{p(e|h_j)}.$$

For a YN task, a decision scheme in our case consists of a statement of the form:

For some value of criterion β and some evidence e, if $l_{12}(e) > \beta$, opt for hypothesis h_1, otherwise opt for h_2 .

On the other hand, for a 2AFC task, the decision scheme could be stated as:

Given evidence e_1 from interval 1 and e_2 from interval 2, if $l_{S,N}(e1) > l_{S,N}(e2)$, opt for interval 1, otherwise opt for interval 2

.

Performance in the YN task depends not only on the perceptual representations of the probability distributions associated with the possible states of the world, but also on the decision criterion β. This quantity is associated with the specific requirement of the YN decision strategy, and is not related to the perceptual representation of the stimuli as such. Performance in the 2AFC task, on the other hand, does not depend on this criterion. We focus on two measures of performance: the percentage of correct responses (PC) and detectability. Detectability is an index of separation between the two internal representations of the states of the world. The hallmark of SR is a non-monotonic behaviour of these quantities as a function of noise intensity, with a local maximum for some optimal amount of noise.

Even though stochastic resonance is reported widely in the literature, for systems ranging from single cells to psychophysical experiments in humans, there is still

uncertainty as to whether the observed effects in humans are genuine or epiphenomenal. The presence of SR in human perception is particularly puzzling because, as stated above, it is currently believed that sensory behaviour conforms to the principles of SDT: from a statistical standpoint, the addition of noise should always result in reduced stimulus discriminability. Furthermore, the extrapolation from single neurons to behaviour is not transparent. A simple simulation reported in this thesis illustrates how SR effects may vanish when averaging across a population of identical neuronal elements, even though the effects are present at the level of each individual element. This example is meant to illustrate how SR effects at the single-cell level may not translate to the level of the neuronal populations that underlie behaviour.
Another potential confounding factor may be sub-optimal placement of the decision criterion β. Under this scenario, spurious SR effects may be measurable in the form of percent-correct values from YN tasks, but not from 2AFC tasks. A result of this kind would indicate that SR does not emerge at the level of sensory perception, but instead reflects sub-optimality of the decision-making process involved in YN tasks.
An additional factor that may complicate the interpretation of SR-like measurements in human behaviour is represented by spatial and/or temporal uncertainty with respect to the location and instant of stimulus presentation. This uncertainty would prompt human observers to monitor detectors over the entire spatial and/or temporal interval over which they expect the stimulus to appear, and such interval may be substantially wider/longer than the actual interval occupied by the signal. This strategy is highly inefficient because detectors that fall outside the spatio-temporal interval occupied by the signal do not provide useful information for performing the task; rather, they degrade performance by contributing their internal noise. This kind of sub-optimality may, under certain conditions, produce effects that masquerade as SR.

We propose an experimental setup designed to minimize the effects of spatial and temporal uncertainty. Furthermore, we combine YN and 2AFC tasks within a unified design that allows direct comparison of the resulting measurements. We collected data using three protocols: one where only YN trials are presented, one where only 2AFC trials are presented, and one where YN and 2AFC trials are mixed and presented randomly on each trial. This approach allows us to test the possibility that SR may reflect sub-optimal placement of the decision criterion: under this scenario, we expect to measure SR effects for PC only on YN trials and not on 2AFC trials. If, on the other hand, SR effects persist for both conditions, we may entertain the notion that SR occurs before the decision-making process and reflects properties of the perceptual representation.

Our pilot results obtained using the YN procedure are suggestive of SR in the human visual system, however some aspects of our dataset indicate the presence of additional factors. For example, we observe non-monotonic behaviour for high noise levels, which is difficult to reconcile with an explanation based on SR. Furthermore, pilot data from mixed-procedure-type trials expose inconsistencies with data from single-procedure trials, raising concerns about the general validity of measurements from the YN procedure: if SR is indeed a characteristic of human sensory perception, it should emerge regardless of the particular modality of behavioural readout. Follow-up studies will need to clarify why SR emerges in pure YN trials but vanishes when trial types are mixed. We will address this issue if the above-noted differences survive additional data collection in a larger subject cohort.

Our immediate concern at this stage is to verify whether our pilot results are robust or due to statistical fluctuation and error measurement. To address this issue, we have collected a larger data sample under controlled stimulus conditions in 10 participants and are developing appropriate analytical tools for determining statistical bounds on our empirical quantities, significance of observed differences, applicability of certain modelling frameworks, and possibly other applications that may become necessary as we inspect our dataset more closely.

Acknowledgements

I wish to thank my supervisor Peter Neri for the help and support in the conceptualization and writing of this thesis. His supervision and input were invaluable, and his infinite patience and understanding made me want to do more.

Furthermore, I wish to thank professors Alessandro Pelizzola and Alfredo Braunstein for the innumerable occasions in which I asked for their advice or sought of them to help me suffer the slings and arrows of outrageous bureaucracy.

Finally, I wish to thank my friends and family. Without them this work would never have seen the light of day. Their support in crucial moments helped me overcome obstacles, be them emotional, situational or academic.

Special thanks to Angelo, Mattia and Denis - the people that stood by me in thought and spirit throughout this journey.

Thank you.

"Who was it who said: 'A man never rises higher than when he does not know whither his path can still lead him.'?"
- Friedrich Nietzsche, *Schopenhauer as Educator*

Table of Contents

List of Figures

Chapter 1

Introduction

Stochastic resonance (SR) is a non-linear phenomenon whereby a subthreshold signal is rendered suprathreshold by the addition of small amounts of uncorrelated noise [1]. Noise components within the energy band occupied by the subthreshold signal push its intensity above threshold; consequently, although in general one may expect signal detectability to decrease monotonically with decreasing signal-to-noise ratio, in this case detectability presents a local maximum for a particular non-zero value of noise. For larger values of noise, the signal is drowned out and the detectability decreases monotonically.

SR is a well known physical phenomenon in systems with non-linear signal detectors, e.g. detectors discriminating via a step function [2][3]. Recent studies have shown that individual cells (or small populations) may exhibit SR effects [4] [5] . These studies have been followed by attempts to demonstrate SR in more complex systems, such as the human visual [6] and tactile sensory system [7], however the evidence remains inconclusive: while some authors argue that SR genuinely occurs at the level of human perceptual discrimination, others maintain that it is an epiphenomenon [8] emerging for reasons such as decision uncertainty [9], suboptimal decision threshold placement [10], non-controlled experimental parameters [11], and possibly others.

Signal Detection Theory (SDT) is an established framework in psychophysics [12]. One of its central goals is to understand the process through which an agent decides whether a particular sensory experience is caused by a stimulus consisting of a signal superimposed on a noisy background, or by noise alone. Thanks to its statistical foundation, SDT represents a powerful theoretical framework for the study of this phenomenon, providing theoretical underpinnings to the experimental procedures widely employed in psychophysical experiments. By and large, it is the theory that has been leveraged in psychophysical work on SR [2].

This report starts with a brief introduction to SDT and an overview of the state-of-the-art regarding SR in complex biological systems, followed by a detailed description of the specific project under way. We focus on those aspects of SDT and SR that are directly relevant to psychophysics. The detailed description of the experimental framework, laboratory setup and data analysis methods refer to what has already been implemented and tested; whenever reference is made to items that are under development and not yet completed, this will be made explicit to readers. The final part of this report puts forward tentative hypotheses for future directions of research; these may be or may not be undertaken, depending upon different experimental outcomes expected in the process of verifying the validity/applicability of our pilot results.

1.1 Signal Detection Theory

The following section is mostly based on "Signal Detection Theory and Psychophysics" by D.M. Green and J.A. Swets [12], a work of fundamental importance in the field, complemented by the review article "Signal Detection Theory, Detectability and Stochastic Resonance" by J. Tougaard [13]. We only introduce the minimal theoretical concepts that are necessary to understand the experimental setup, stimulus design and our choices of data analysis, without attempting a full account of SDT or SR (such an account would far exceed the remit of this report).

SDT is foundational to contemporary sensory psychophysics. In a classical sense, psychophysics studies and models the "relationship between stimulus and sensation" [14]: it seeks to explain the connection between objective physical stimulation and the subjective perception/sensation of a biological organism in response to said stimulation. In the modern sense of contemporary formulations, the above connection is statistical, hence the need for a theoretical framework rooted in statistical decision theory.

Practically speaking (i.e. during experimentation in the laboratory), the minimal requirements for a psychophysical measurement are the following [12]:

1. A **task** for the subject to perform. The specification of the task determines the framework through which the subject will perceive stimuli and reach decisions. Furthermore, the setting of the task allows for the classification and evaluation of the procedure used to reach a decision
2. At least 2 possible **states of the world** represented by physical stimuli, most often construed as pure noise (n) or a signal superimposed on a noisy background (s)
3. One or more temporal/spatial **presentation interval(s)** during which the subject is provided with sensory information regarding the state of the world
4. The subject's **response**, generated by some sort of behavioural output (e.g. button press), reflecting a decision about which state of the world most probably generated the stimulus associated with that response

In psychophysical experiments a subject is presented with an input which they then categorize as being caused by one of the possible states of the world, and their response is registered. The possible responses are often, but not necessarily, in a 1-to-1 correspondence with the states of the world. Put in other words, the subject's goal during an atom of the psychophysical experiment is to detect a potentially presented signal.

As already mentioned, the ideas of Signal Detection Theory are formalized through Statistical Decision Theory, which studies how an agent reaches optimal decisions when presented with an input characterized by uncertainty [15]. In the following we recall some cardinal concepts and results that will prove useful later.

1.1.1 General concepts

Any well-specified behavioural decision is associated with an equally well-specified goal. Not only does having a goal give a particular decision scheme its meaning, but it also makes it possible to choose some decision scheme to begin with. For protocols that fall directly under the SDT framework, the goal amounts to giving as many correct answers as possible in a given task. Because the input stimulus is characterized by some degree of uncertainty (typically in the form of stimulus noise), it is often the case that no decision strategy can achieve perfect performance (correct response on every stimulus instance); however, we can say that, given a specified objective, a particular decision rule works better than others - *on average* (i.e. in a statistical sense).

We define each instance of a stimulus-response event as 'trial'. On every trial, either an s or n stimulus is presented, and the subjects' goal is to determine whether the observed stimulus originated from the s or the n configuration by responding to the question "*Was a signal presented in the last trial?*". They would be correct [incorrect] to answer '*yes*' ['*no*'] if s was presented, and '*no*' ['*yes*'] otherwise.

There are 3 fundamental elements to any decision, of which we consider the discrete versions:

- $h_j \in \{h_j\}_{j=1}^{N}$, the actual state of the world, usually characterized by presence or absence of a signal
- $e_k \in \{e_k\}_{k=1}^{M}$, the information provided to the subject regarding the state of the world; this can also be thought of as evidence supporting different states to different degrees and it can be an either qualitative or quantitative item, as well as a combination of multiple elements of both types
- $H_i \in \{H_i\}_{i=1}^{N}$, the subject's decision regarding what the state of the world is, given the evidence at hand

Subjects are always instructed on the possible stimuli they may witness, and the corresponding states of the world that generate them. In our case, we restrict ourselves to 2 possible states of the world: signal+noise (s) and noise alone (n). From the viewpoint of SDT, the perceptual system constructs an internal representation of the probability distributions related to the possible states of the world. These, combined with a decision criterion, support a behavioural decision. The formalism of Bayesian statistics lends itself rather well here. We call *prior* the probability that the j^{th} state of the world is the actual state, and denote it $p(h_j)$. Secondly, the posterior is the probability that the state of the world i is the real one, given that some evidence has been observed, mathematically expressed as $p(h_i|e_k)$. Thirdly, the likelihood is written as $p(e_k|h_i)$, and is the probability that a piece of evidence e_k is supportive of the state of the world h_i. We can use likelihood to compute the *likelihood ratio*:

$$l_{ij}(e_k) = \frac{p(e_k|h_i)}{p(e_k|h_j)} \tag{1.1}$$

The likelihood ratio is a measure of the legitimacy attributed to hypothesis h_i by evidence e_k. $l_{ij}(e_k) \in (0, 1)$ favours hypothesis h_j, whereas $l_{ij}(e_k) \in (1, +\infty)$ favours h_i.

In general, a decision scheme consists of a statement of the form:

For some value of criterion β and some evidence e_k,
if $l_{ij}(e_k) > \beta$, opt for hypothesis h_i, otherwise opt for h_j.

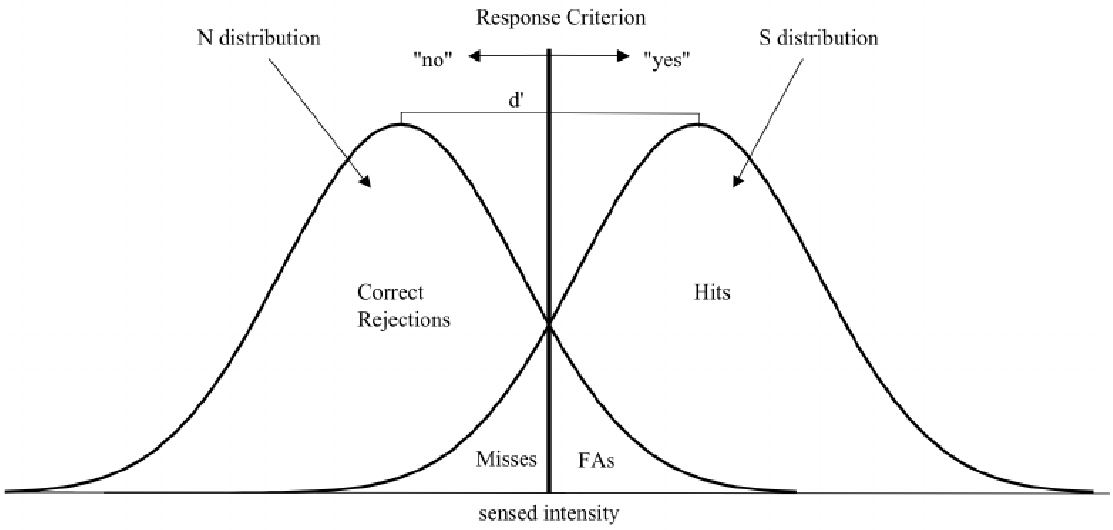

Figure 1.1: S and N probability distributions, the decision criterion, and the detectability d' (adapted from [16])

A critical underlying assumption is that subjects are capable of mapping arbitrary stimuli to an internal metric scale, and hence indirectly to one another. This comparison is then used to decide whether a given stimulus is associated with a given state of the world, e.g. s or n. It is not necessarily the case that subjects weigh evidence as prescribed by the likelihood ratio: they may adopt different rules. However, it can be proven that, given a criterion β, the likelihood ratio is the optimal decision variable [12], as is any other decision variable that is an increasing monotonic function of the likelihood ratio.

As mentioned earlier, for a given presentation of evidence regarding the true state of the world, SDT makes the reasonable assumption that subjects can be modelled as sampling from one of two probability distributions, depending on the true state of the world. For two distributions with the same variance, simple optimization shows that the optimal decision criterion lies exactly between the means of the distributions. It is often assumed that the distributions of s and n are Gaussian distributions, characterized by the same variance σ, with their respective means μ_s and μ_n shifted from one another. Fig.1.1 features both distributions with the optimal decision criterion.

The critical feature that characterizes the state of affairs depicted in Fig.1.1 is the degree to which the two distributions are distinguishable from one another. Ultimately, this is the property that will affect any decision based on those distributions, and this property can be quantified via the detectability index d':

$$d' = \frac{\mu_s - \mu_n}{\sigma} \tag{1.2}$$

Eq.1.2 specifies a unit-less metric for quantifying the separation between the two distributions associated to the evidence for the different possible states of the world: for a given spread (unreliability) of the two distributions, increasing the distance between them will produce larger d' values; similarly, for a given distance between the two distributions, reducing spread will increase d'. It is important to notice that d' does not depend on the criterion β since the latter does not enter Eq.1.2. β is merely a way to convert information from the two distributions onto a decision rule of the kind exemplified above. Another important remark is that, because σ is monotonically related to the amount of external noise injected into the stimulus (more stimulus noise $\rightarrow$ greater σ), d' is expected to decrease monotonically with increasing external noise. As mentioned briefly before, SR involves a non-monotonic relationship between the discrimination performance of a sensory system and the amount of external noise applied to the stimulus; if performance is assessed with reference to d', we must conclude that SDT does not predict SR.

In order to verify whether the above prediction experimentally applies or not, in the laboratory d' is estimated using various procedures. For the purposes of this report, we focus on the Yes-No task (YN) and the 2 Alternative Forced Choice task (2AFC). Even though the type of stimuli presented to the subject remain the same in both tasks, the information given to the subject and the subsequent responses differ. Furthermore, because of these differences, each procedure has its own set of advantages and shortcomings, which must be leveraged or controlled during an experiment. More specifically, as we discuss in more detail below, the difference between these two procedures lies in their ability to support an empirical estimate of d' that does not depend on response criterion β.

1.1.2 Yes/No Procedure

In a Yes/No procedure, every trial generally consists of a warning stage, a presentation stage, a decision stage and a feedback stage. These usually occur as a function of time. The objective is usually to give as many correct answers as possible.

The warning stage consists of a cue that alerts the subject that a stimulus is to be presented shortly. Cues can be given through any sensory modality; visual and auditory cues are the most common ones. Even though this stage is somewhat unnatural, in the sense that warning cues almost never happen in real life, it is often introduced to alert subjects and maintain vigilance.

During the presentation stage, subjects are presented with a stimulus featuring *either* signal+noise *or* noise alone. The typical question is "Did the last presentation of the stimulus feature the signal?", and the subject is only allowed to answer either "Yes" or "No". In some cases, feedback is provided to the subject, informing them whether their decision was correct or not. It is often assumed that decisions from any two separate trials can be treated as statistically independent, although some degree of inter-trial dependence is known to exist and has been extensively characterized [17].

After having been exposed to a stimulus, the subject's response can then be either S or N. We use the convention of writing in lowercase the real states of the world, and in uppercase the subject's responses. This yields 4 possible outcomes of a decision event:

- $S|s$, called a *hit*, denoted by H
- $N|s$, called a *miss*, denoted by M
- $S|n$, called a *false alarm*, denoted by FA
- $N|n$, called a *correct rejection*, denoted by CR

After N trials, each of the four possible outcomes will make up a fraction of the total. Each fraction is an estimate of the conditional probability that the subject may opt for a hypothesis, given what the actual state of the world is. For each state of the world, the probabilities of a decision conditional upon that state of the world must add to one. For example, the probability of a hit and the probability of a miss must sum unity. For this reason, there are only two degrees of freedom for the four possible outcomes: all relevant information about the average behaviour of the stimulus-response coupling can be captured, for example, by $p(S|n)$ and $p(S|s)$.

		Response Alternative: S	Response Alternative: N
Stimulus Alternative	s	hit $p(S\|s)$	miss $p(N\|s)$
	n	false alarm $p(S\|n)$	correct rejection $p(N\|n)$

Table 1.1: Stimulus Response Matrix

How do we assess 'performance' in a yes-no task? The most straightforward approach would be to calculate the percentage of correct responses (PC), i.e. $(\hat{p}(S|s)+\hat{p}(N|n))/2$ where $\hat{p}$ is the estimated probability from data (average of trials of a given type). There is a potential confound with this measure of performance: even under the simplest SDT model (Fig.1.1), it depends on the criterion β. Imagine setting β to the extreme right of the plot in Fig.1.1: under this scenario, $\hat{p}(N|s)) = 1$ and $\hat{p}(S|s)) = 0$ (the observer always responds "no"), and percent correct is 50% i.e. chance. Based on PC, we would conclude that the observer's internal representation of the stimuli carries no discriminatory power: it produces a response that is equivalent to pressing buttons randomly. However, this conclusion would be incorrect: the internal representation is in principle capable of discriminating between the two distributions associated with the two states of the world in Fig.1.1; the problem is not with the internal representation, but with the manner in which the observer has placed their criterion β when reading off that internal representation for the purpose of producing a binary yes/no response. In other words, it is entirely possible that the observer is able to perceive the difference between stimuli generated by state-of-the-world s and stimuli generated by state-of-the-world n, but applies a response criterion that does not reflect this ability. When the criterion is placed in such sub-optimal way, we speak of 'response bias'.

For a different way of expressing the concept discussed above, consider that response bias can be influenced by prompting subjects to either avoid one type of error, or prefer a certain type of correct response. For instance, a subject may become less inclined to answer S if the punishment associated with a false alarm is increased. On the other hand, propensity towards answering N may be increased by increasing the reward for correct rejections, leading to a decrease of S responses. Through such manipulations of the goal, one may influence a subject's response bias. We do not suppose, however, that the subject's perceptual representation is affected by this class of manipulations: in the most parsimonious account of sensory processing, stimuli are perceptually represented with a given fidelity, regardless of how the associated behavioural responses are rewarded/punished. In other words, the issue of how one goes about *utilizing* the information contained within the perceptual representation for avoiding/prioritizing specific outcome behaviours is viewed as separate from the issue of how much discriminatory power is carried by the perceptual representation itself. Because percent correct depends on the former (see above), it is not an appropriate metric for estimating the latter.

If we assume that the perceptual process is well-approximated by the scenario depicted in Fig.1.1, then we *can* obtain an estimate of discriminatory power that is independent of response criterion, i.e. a more direct estimate of d' itself, via the following expression:

$$\hat{d'} = \Theta^{-1}(\hat{p}(S|s)) - \Theta^{-1}(\hat{p}(S|n)) \tag{1.3}$$

where Θ is the cumulative distribution function of the standard normal distribution. It must be emphasized that this expression is not, in general, an unbiased estimate of d' (as defined in Eq.1.2); in order for it to be so, we must assume that the underlying sensory process conforms to the specifications of Fig.1.1: two Gaussian distributions of equal variance subjected to a response criterion threshold.

1.1.3 2 Alternative Forced Choice Procedure

Another type of procedure is the 2 Alternative Forced Choice procedure. Generally speaking, the goal assigned to subjects remains the same - responding correctly to as many trials as possible. All stages of this procedure are exactly the same as for the Y/N procedure, with only the presentation and the response stage differing noticeably. In this procedure, both states of the world are shown on every trial. The presentation of the two states can occur as a temporal sequence, presenting the two states one after the other, or simultaneously at different spatial locations. The subject must then decide which of the two presentations most likely contained the signal. A general decision scheme for this type of setting consists of a statement of the form:

Given evidence e_1 from interval 1 and e_2 from interval 2,
if $l_{S,N}(e_1) > l_{S,N}(e_2)$, opt for interval 1, otherwise opt for interval 2.

As opposed to the decision rule introduced earlier, the rule immediately above does *not* involve any response criterion β: the likelihood that the stimulus in the first interval contains s as opposed to n ($l_{S,N}(e_1)$) is directly compared to the likelihood that the stimulus in the second interval contains s as opposed to n ($l_{S,N}(e_2)$); whichever interval returns the highest likelihood is selected as containing the target stimulus.

If we assume that likelihoods are computed from the underlying sensory representation depicted in Fig.1.1, we must conclude that the PC is monotonically related to d': in this case, there is no issue with the potential role of criterion bias as encountered when discussing the YN procedure (section 1.1.2), because the notion of a criterion does not enter the decision rule. We are therefore left with a *criterion-free empirical estimate of performance* that is transparently related to d', and therefore to the discriminatory power of the underlying sensory representation.

It is important to notice that the 2AFC decision rule contains an unstated assumption: subjects are not expected to demonstrate any intrinsic bias towards selecting either interval. In the example of spatial sequencing with the stimuli appearing at two different spatial locations, the subject is assumed to have no *a priori* preference for perceiving the signal appearing on the right, for instance, as opposed to a signal appearing on the left (under conditions when the probability of these two events is equal). Similarly, when stimuli are presented in temporal sequence, subjects are assumed to have no preference for either the first or the second interval. Although this assumption seems reasonable, it does not necessarily apply and in fact does not apply under certain conditions (particularly with temporal sequencing), making it necessary to verify its applicability from data.

1.2 Stochastic Resonance

Stochastic resonance is a statistical, non-linear phenomenon whereby the detectability of a signal can be positively affected by the addition of relatively small amounts of noise [1][2] [18] [19] [20]. The phenomenon is usually discussed in the context of subthreshold signals. In general, for a signal that is slightly below detection threshold at all times, there exists some optimal amount of noise for which the detectability reaches a maximum, before decreasing monotonically.

This phenomenon is typically demonstrated for a sinusoidal signal whose intensity is slightly below the detection threshold of a detector (Fig.1.2). Subsequently, a zero-mean gaussian noise is added to the signal. The noise need not necessarily be white, however its upper cut-off frequency is usually much higher that that of the signal. A consequence of this superposition is that some of the noise frequencies resonate with the appropriate signal frequencies, pushing them above the threshold level. Whenever they upwardly surpass the threshold, the detector fires. Subsequently, a random sequence of spikes can be associated with detector activation, and the average instantaneous firing rate is modulated by the amplitude variation of the signal. This translates to the power spectrum of the detected signal so that there is a spike associated with the resonant frequency, surrounded by a relatively flatter or vanishing power spectrum.

The detectability of a signal, when expressed in terms of the signal-to-noise ratio reads

$$d' = \frac{S_{s+n} - S_n}{S_n} \tag{1.4}$$

where S_{s+n} and S_n are the spectra of the detected signal and the noise, respectively. Generally, it is expected that the detectability of a signal decreases with the SNR. However, a hallmark of SR is precisely a local maximum in the detectability of the signal at some non-zero SNR. As the noise increases, the signal becomes overwhelmed by it, so the detectability is a concave function of SNR [20].

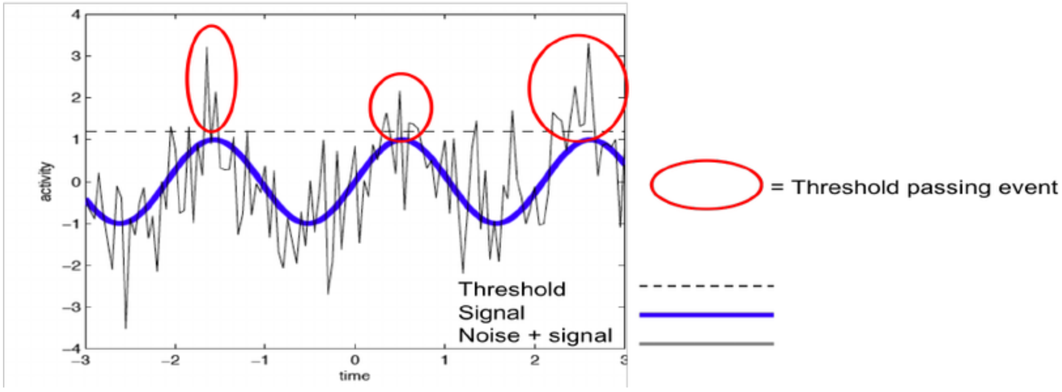

Figure 1.2: Paradigmatic example of Stochastic Resonance (adapted from [21])

In the context of perceptual discrimination, SR is defined as a variation of detection performance that is non-monotonic with increasing stimulus noise (decreasing stimulus SNR), reaching a maximum for a non-zero value of stimulus noise. The phenomenon of SR is interesting from a neuroscientific standpoint due to the inherent non-linearity of cortical structure and the ubiquity of noise, both external and internal, in all sensory systems. It is therefore interesting from the viewpoint of sensory perception and information processing. While some authors claim to have demonstrated the presence of SR in human sensory discrimination [6], others are of the opinion that it merely represents an epiphenomenon of possibly artifactual origin [8]. As further clarified below, a major drive behind this study was to settle this controversy conclusively.

Chapter 2

Motivation

Even though the effect of SR within the human central nervous system superficially may seem of relatively minute import, the reality may be substantially different. The first part of this section outlines the wider scope of the study of SR in humans, in an attempt to show the reader how the study of this phenomenon may contribute to the understanding of the human experience. The second part of this section provides motivation related to this particular project, and serves to elucidate doubts and problems that have been (and are still being) considered in its ideation and implementation.

In psychophysical experiments decisions are communicated via motor outputs, or in other words, sensory phenomena are filtered through the nervous system and subsequently result in behavioral output. Suppose that a subject hypothetically perceived some stimulus without having been given any objective (so no response mediated by behavior is required of him), and suppose it were possible to study the nature of the subject's internal representation of the event as informed by their perceptive systems without them communicating their decision. One question that arises is whether the perceptual system provides the individual with a definite or a statistical representation of the world?

Drawing on experiences from normal, conscious existence one may be inclined to state that the representation of the phenomena perceived in the world are of a definite nature: they are either perceived or are not perceived. A counterargument to that position is that such binary, definite representation is a byproduct of a preponderance of intense enough evidence pointing to the presence (or absence) of what is being witnessed (or not). If, however, the evidence were characterized by a great deal of inherent uncertainty, the perceptive system may "prefer" generating an internal representation which is statistical in nature, only to be rendered definite by the decision process. Stated differently, it may be the case that during the perception phase different events are assigned probabilities through some internal process, only to be categorized in a definitive manner via the application of a

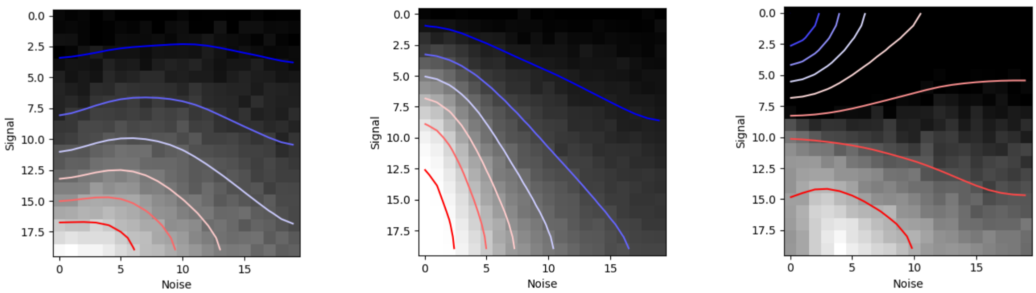

(a) (b) (c)

Figure 2.1: Plotted performances of simulated neuronal elements: Fig.2.1(a): Signal detection in single neuron; Fig.2.1(b): Signal detection in neuron population (2AFC); Fig.2.1(c): Signal detection in neuron population (YN). Surface brightness reflects detection performance (brighter for better performance).

decision criterion.

SR is inherently connected to the above question. A simplified instance of this dichotomy is the difference in output between a single neuronal element and a population thereof, illustrating how an SR effect at a single cell level need not necessarily translate to the level of complex systems. To that end, Fig.2.1 contains plots obtained from a computer simulation activations of single and pooled neuronal elements when given an input characterized by varying degrees of noise. The neuronal element only fires if the presented input is higher than an activation threshold.

Fig.2.1(a) shows the PC of a simulated single neuronal element endowed with a threshold and a low level of internal noise relative to the external noise intensity. SR can clearly be seen, especially for medium intensity signals. To understand how SR is visible at the level of the surface plot in Fig.2.1(a), imagine taking a horizontal slice across this plot for a signal level (y axis) that is about 1/3 away from the origin along the y axis: if you were to plot this slice as a function of noise intensity (x axis), it would look non-monotonic, peaking for a specific non-zero value of noise intensity. Fig.2.1(b) shows performance obtained by pooling 100 simulated neuronal elements sharing the same threshold while differing in internal noise, and averaging their responses over 5000 2AFC type trials. It is immediately obvious that all traces of SR vanish: any horizontal slice across the surface plot will be monotonically decreasing, with maximum performance corresponding to 0 noise intensity. When the 100 neurons are exposed to an input, each randomly classifies the input depending on its instantaneous internal state. When averaging across these multiple noisy decisions, the original statistical structure of the underlying response distributions (before thresholding) is restored (albeit with loss of information). As predicted by SDT, this structure should not manifest SR (see section 1.1.1).

Referring back to the discussion of definite vs. statistical representation, this simulation serves as an illustration of how a statistical representation may emerge from an amalgamation of lower level definite (binary) representations. The statistical representation can be again converted to a definite state through a *post hoc* application of a decision criterion. However, it is important to highlight the subtle distinction that, whereas in one case perceived items get definitively represented at the perceptual level, in the other they are statistically represented at the perceptual level, only to get definitively represented at the behavioral level via a decision process. It is conceivable that the study of SR viewed through this lens could contribute to the understanding of information processing and representation within the human cortex, as well as to what happens on the interface between unconscious and conscious perception.

In recent years, SDT has also been brought to bear on the problem of SR in the human brain. The problem of decision criterion emerges in that context. In particular, the decision criterion during YN trials is usually sub-optimally placed due to the subject's idiosyncrasies, which can lead to PC showing apparent effects of SR. If these effects are epiphenomena, they should disappear when the same data is used to compute d' since the latter by definition 1.2 represents the discriminability of stimuli at the perceptive level, and is independent of behavior.

Another way of controlling for the bias in YN trials is by comparing the obtained results with those from 2AFC trials, since they are completely unbiased if the requirements from section 1.1.3 are met. If SR effects are a secondary phenomenon caused by sub-optimal bias placement, then they should not appear in the PC obtained from 2AFC trials. Incidentally, this is exactly what can be seen by comparing Fig.2.1(b) and Fig.2.1(c). A way for controlling for this is to either implement YN type procedures and study the signal detectability, or to implement a 2AFC procedure only.

A final concern relates to spatial and temporal uncertainty [9][11]. When a subject is uncertain as to when or where a stimulus may appear, they apply a bank of multiple detectors, each dedicated to a portion of the spatial and temporal field. The input information is then sent through channels associated to those detectors and is integrated across the whole spatial and temporal range. The result is not efficient because not only does each detector have its own internal noise, but the input for each detector potentially differs. If, for instance, a signal appears on only a fraction of the total spatial range, only the associated detector may get a signal input, while the others read out only noise. This sub-optimality may, under certain conditions, produce SR-like effects [11]. A way to control for this is to reduce as much as possible the uncertainty related to the stimulus presentation. This can be done through careful addition of cues and markers to the experimental setup.

Chapter 3

Methods

3.1 Experimental setup

The central element of the experimental setup consists in a computer code written in *Python3* [22], integrating elements of the open source package *PsychoPy*[23]. It encodes for each aspect of the decision cycle and the data collection.

Healthy participants are placed in front of a screen, at a distance of approximately $57cm$, so that a stimulus of height and length equal to $1cm$ occupies approximately 1 degree of visual angle. They are instructed on the nature of the task, and asked to give as many correct answers as possible. No incentives are given to either avoid or prefer any particular type of outcome. Each sequence of trials contains between 600 and 2000 trials, depending on a participants' stamina and aptitude. We expect to collect a total of $10000 - 12000$ trial responses per subject. Each subject sits through one or more sequences, potentially over a span of multiple days or weeks.

An initial phase consists in participants being exposed to sequences containing solely "blocks" of Yes/No procedure type trials in order to confirm that our experimental setup is indeed capable of detecting SR effects. Once that is confirmed, the second phase consists in sequences containing the same number of randomly mixed YN and 2AFC trials. One important aspect of the mixed trial type paradigm is that it makes it impossible for the subject to adapt to a certain type of stimulus layout. This adaptation could in turn bias the decision process and introduce fictitious SR effects. Another important aspect is that it permits data collection of two different procedures roughly in parallel, as opposed to sequentially, in a span of hours or days. It can then be reasonably assumed that the resulting data set is characterized by a uniform degree of internal noise, allowing a more direct comparison between results of different procedure type trials.

3.1.1 Stimuli

Every stimulus item is composed of 13 vertical, monochromatic bars. The length and width of the bars is such that the stimuli are quadratic and subtending 2 degrees of visual angle. Each bar's shade of grey is drawn from a 13-by-13 stimulus matrix. A matrix row contains 13 different, gaussian-distributed values, whereas each matrix column contains one value only, as can be understood from Fig.3.1. The grey-scale of the monitor ranges from black to white, quantified respectively by 0 and 1. For a healthy human subject, a barely but reliably detectable pulse is expected to lie in a range of contrasts between 0.003 and 0.015 relative to the grey background at 0.5.

The shades of grey of the noise are distributed according to a gaussian distribution centered in 0.5, and with a variance equal to a percentage of the intensity of the pulse drawn randomly for every trial from the set (0%, 3%, 5%, 10%, 20%, 40%). The particular value of the threshold pulse intensity is preliminarily estimated for every subject separately and an array of $2-3$ values of intensity distributed around that threshold intensity is chosen. The pulse intensity will randomly be sampled from that array at each trial. As reported in Fig.3.1, the difference between an n and an s state is that the latter features the pulse added to the central (7^{th}) bar. The presentation of a stimulus can occur on either the left of the right half of the screen. Regardless of procedure and state, the geometrical center of a stimulus always has an eccentricity of 2.5 degrees of visual angle with respect to the geometrical center of the screen along the x axis, and no eccentricity with respect to the geometrical center of the screen along the y axis.

3.1.2 Procedures and sequences

At the start of every trial, the screen features a white fixation cross which the subject is instructed to fixate, and four white bars. The purpose of the bars is to mark the locale of the pulse presentation, thus reducing spatial uncertainty. These markers are static. The warning stage starts $150ms$ before the presentation stage, and they end simultaneously. In the mixed procedure modality its purpose is threefold: firstly, it reduces the temporal uncertainty the subject may otherwise experience regarding the onset and duration of the presentation stage; secondly, it removes uncertainty related to whether the forthcoming presentation is of YN or 2AFC type; thirdly, in YN type procedures it indicates on which side of the screen will the stimulus appear. Fig.3.2 shows the presentation stage midway, with all the elements in full view. Fig.3.1(a) features a YN trial with a pulse appearing on the left-hand side of the screen, and Fig.3.1(b) shows a 2AFC trial with signal+noise on the left, and just noise on the right-hand side of the screen. The presentation stage lasts $250ms$, after which the orange cues disappear and the fixation cross turns white again.

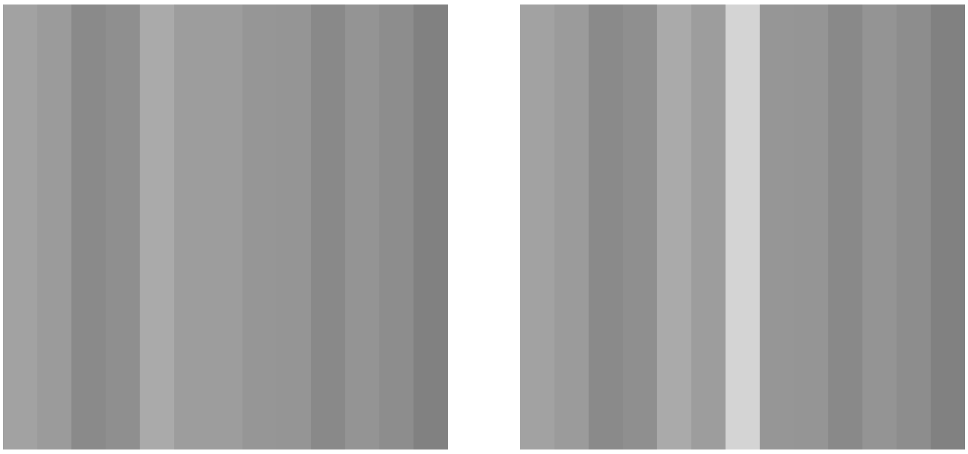

(a) Stimulus with pure noise (b) Stimulus with pulse + noise

Figure 3.1: Example of a pure stimulus (a) and a stimulus with a pulse added to the noise (b). Both the pulse intensity and the noise variance have been adjusted in order to make the stimuli more readily visible.

In the YN trial the subject communicates their answer to the question "*Was there a signal* presented during the last trial?" by pressing either y (for "yes") or n (for "no"). Alternatively, in the 2AFC trial they respond to the question "*Where was the signal* presented during the last trial?" by pressing either the left of right arrow key. Keys "y" and "n" are not accepted as answers to 2AFC trials, and vice versa. After the participant has given an answer, the fixation cross turns red or green depending on whether the response was incorrect or correct, respectively. The resetting of the fixation cross' color to white marks the end of the trial. Every 50 trials the participant is informed of their performance by displaying the PC they have given during the sequence up to that point.

3.2 Data analysis

The set of tools for data analysis are incomplete since the experiment is still in the pilot stage. What will mostly be considered is the PC for 2AFC trials, and the d' for YN trials. Plots of those quantities as functions of external noise level will initially serve as a crude indication of whether SR effects emerge from the collected data.

Another aspect of data analysis may include what we call "Response Conditioned Averages" or RCA, an inverse correlation technique which will potentially give an insight into the perceptual template the subject applies when a stimulus is presented

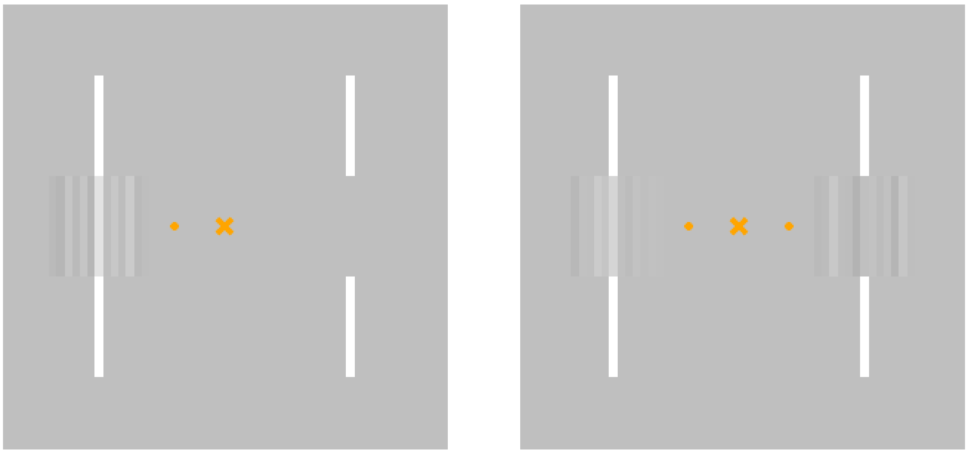

(a) Yes/No presentation (b) 2AFC presentation

Figure 3.2: Example of the presentation stage for the 2 procedures employed in the experiment. As in Fig.3.1, the pulse and the noise intensity have been adjusted in order to make the stimuli more readily visible.

during a YN type trial. In practice, responses to YN trials are placed in 4 categories, depending on the outcome: hits, false alarms, misses, and correct rejections. Within each category, the average is taken over the stimulus matrices. After that the matrices associated to N answers are subtracted from those associated to S answers. The final result is referred to as "meta template" or simply "kernel". Considering the baseline to be at 0.5, we expect the noise values in N response trials to be on average negative enough to mask the pulse, whereas in the S response trials we expect the noise values to be either close to 0 or positive - so they enable the recognition of a pulse or even simulate it. Due to reasons beyond the scope of this report, the values of each row of the resulting matrix should mimic a Mexican hat function or Ricker wavelet. Inspecting the obtained wavelets in function of different levels of noise, given a pulse intensity, may inform us as to whether and how the perceptual template of the visual system changes depending on noise levels in this sort of task.

Chapter 4

Results and Conclusions

4.1 Results

In this section we present results thus far obtained during the pilot stage. These are by no means conclusive, but may provide some legitimacy to our experimental setup and the choices made along the way.

We have mentioned that the PC in YN trials may display SR due to sub-optimal criterion placement, i.e. individual biases. This fictitious effect can be avoided by plotting the detectability d', or performing 2AFC trials. Fig.4.1 shows plots obtained from 4200 pure YN trials for Subject 1 and 600 trials for Subject 2. Note that the noise vales in Fig.4.1(a) are not exactly the same as in the array reported in subsection 3.1.1. However they still represent noise levels as fractions of signal intensity. Both figures indeed show an increase in detectability for a non-zero noise level. Incidentally, for the same signal intensity both subjects display maximum SR effects for the same quantity of noise added, i.e. 10% of signal intensity.

In figures (4.1(c),(d)), the x axis shows the positions of the stimulus bar relative to the central bar, whereas the y axis contains numerical values associated to shades of grey relative to the backdrop at 0.5. Indeed, it can be seen from the figures that the Ricker wavelet shape is emerging. It is noteworthy that the wavelet is much better defined in Fig.4.1(c) than in Fig.4.1(d), with the former being a much larger data set.

Fig.4.2 shows plots from mixed-procedure-type trials for Subjects 1 and 2. It can be seen that the detectabilities again display non-monotonic behavior as noise increases. However, for Subject 1 the maximum is at 3% of signal intensity, instead of 10%. Furthermore, no trials with signal intensity 0.003 and 40% noise were sampled for Subject 2, due to the restricted size of the sample. The plots for PC in the 2AFC trials also give no definitive evidence of SR. On the one hand, Subject 1 exhibits only a minuscule peak in performance at noise equal to 3% of signal

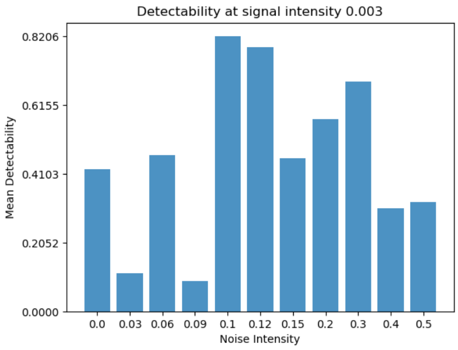

(a) Detectability for Subject 1

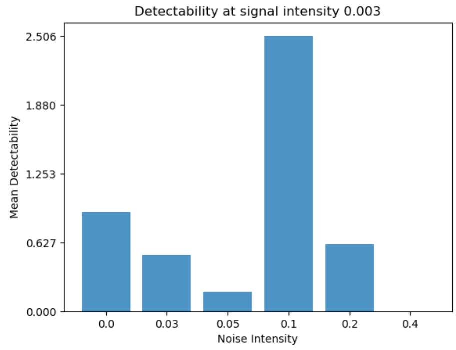

(b) Detectability for Subject 2

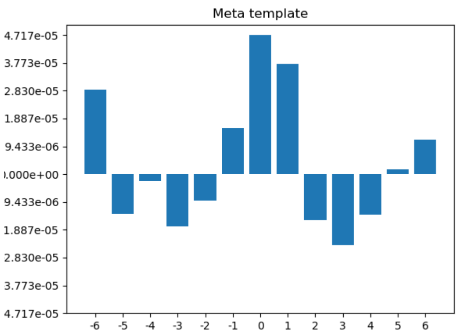

(c) Kernel for Subject 1

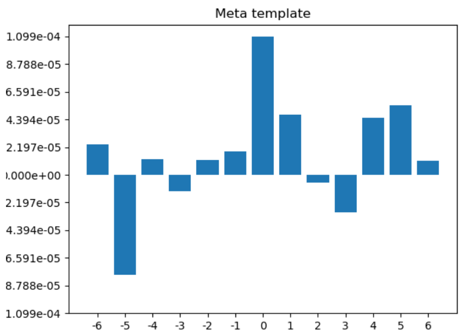

(d) Kernel for Subject 2

Figure 4.1: Plots obtained from pure YN trials

intensity, falling within less than 2 standard deviations of the performances at adjacent values of noise intensity. On the other hand, Subject 2 exhibits roughly the same performance at 5% and 40% noise, indicating that both peaks could be due to statistical fluctuations.

Finally, the perceptual meta-templates, while retaining some of the fundamental features of Ricker wavelets, show no convincing structure.

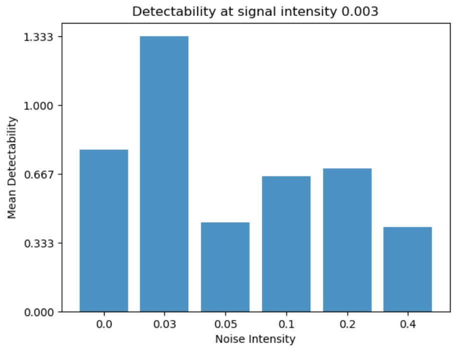

(a) Detectability for Subject 1

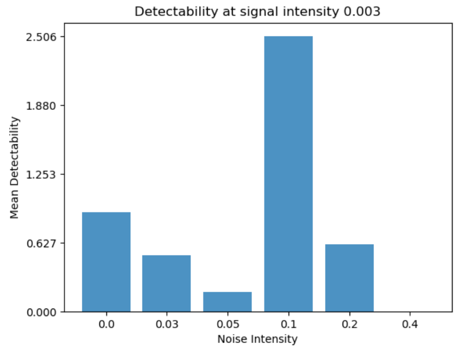

(b) Detectability for Subject 2

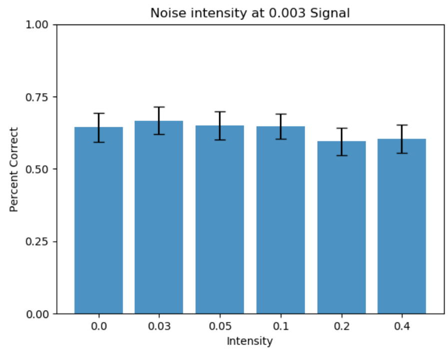

(c) PC in 2AFC for Subject 1

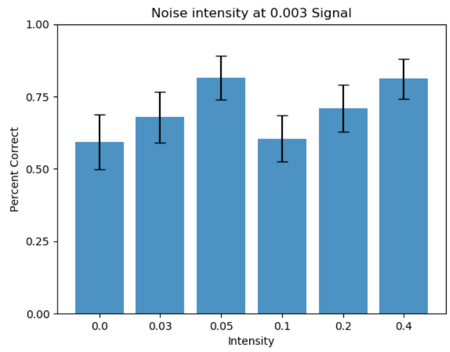

(d) PC in 2AFC for Subject 2

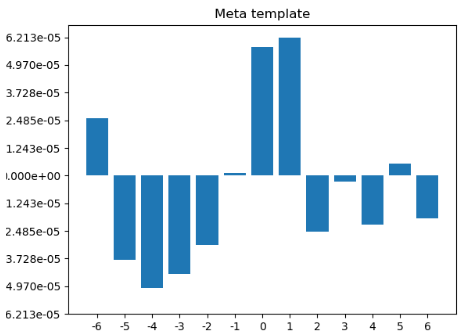

(e) Kernel for Subject 1

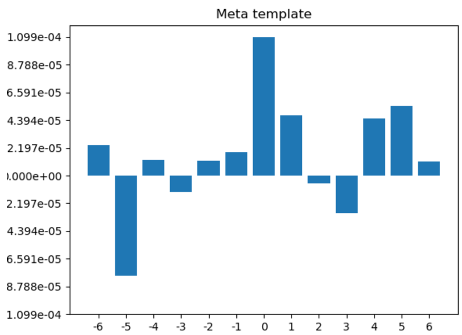

(f) Kernel for Subject 2

Figure 4.2: Plots obtained from mixed YN\2AFC type framework

4.2 Conclusions

The pilot results obtained using the YN procedure are indicative of SR in the human visual system. However, the relatively small sample and the non-monotonicity of the plots for high levels of noise are issues that need to be addressed through more data collection and analysis. Furthermore, a variant of YN trials may be introduced in order to further reduce spatial uncertainty, namely by always presenting the stimulus on only one side of the screen. Because this design may tempt observers to move their eyes away from fixation and directly onto the expected location of stimulus appearance, their eye position will be tracked to ensure that only trials during which subjects looked at the fixation cross are retained for further analysis.

In contrast to single-type YN trials discussed above, pilot data from mixed-procedure-type trials appear to indicate that SR in humans is epiphenomenal - specifically referring to 2AFC trials. It could be argued that, if SR is indeed a characteristic of human sensory perception, it should emerge regardless of the particular modality of behavioral readout. In other words, it should persist regardless of whether the task at hand is of YN or 2AFC type, and regardless of whether the procedure type is pure or mixed. Follow-up studies will need to clarify why SR emerges in pure YN trials, but vanishes when trial types are mixed. We will address this issue if the above-noted differences survive additional data collection in a larger subject cohort (see below).

Our immediate concern at this stage is to verify whether our pilot results are robust or due to statistical fluctuation and error measurement. To solve this issue, we have collected a larger data sample under controlled stimulus conditions in 10 participants, and are developing appropriate analytical tools for determining statistical bounds on our empirical quantities, significance of observed differences, applicability of certain modelling frameworks, and possibly other applications that may become necessary as we inspect our dataset more closely. If the observed discrepancy between YN and 2AFC persists, it will be necessary to understand the underlying perceptual mechanism. We will do this through further testing with targeted stimulus manipulations, as well as with computational modelling whenever appropriate.

Bibliography

[1] R Benzi, A Sutera, and A Vulpiani. «The Mechanism of Stochastic Resonance». In: *Journal of Physics A: Mathematical and General* 14.11 (Nov. 1981), pp. L453–L457. DOI: 10.1088/0305-4470/14/11/006. URL: https://doi.org/10.1088/0305-4470/14/11/006 (cit. on pp. 3, 11).

[2] Frank Moss, Lawrence Ward, and Walter Sannita. «Stochastic Resonance and Sensory Information Processing: a Tutorial and Review of Application». In: *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology* 115 (Mar. 2004), pp. 267–81 (cit. on pp. 3, 11).

[3] François Chapeau-Blondeau. «Stochastic Resonance and the Benefit of Noise in Nonlinear Systems». In: *Noise, Oscillators and Algebraic Randomness*. Ed. by Michel Planat. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 137–155 (cit. on p. 3).

[4] Sergey Bezrukov and Igor Vodyanoy. «Reply: Stochastic resonance at the single-cell level». In: *Nature* 388 (Aug. 1997), pp. 633–633. DOI: 10.1038/41688 (cit. on p. 3).

[5] J. J. Collins, T. Imhoff, and P. Grigg. «Noise-enhanced information transmission in rat SA1 cutaneous mechanoreceptors via aperiodic stochastic resonance.» In: *Journal of neurophysiology* 76 1 (1996), pp. 642–5 (cit. on p. 3).

[6] Keiichi Kitajo, Daichi Nozaki, Lawrence Ward, and Yoshiharu Yamamoto. «Behavioral Stochastic Resonance within the Human Brain». In: *Physical review letters* 90 (June 2003), p. 218103. DOI: 10.1103/PhysRevLett.90.218103 (cit. on pp. 3, 12).

[7] Wen Liu, Lewis Lipsitz, Manuel Montero-Odasso, Jonathan Bean, D. Kerrigan, and James Collins. «Noise-Enhanced Vibrotactile Sensitivity in Older Adults, Patients with Stroke, and Patients with Diabetic Neuropathy». In: *Archives of physical medicine and rehabilitation* 83 (Mar. 2002), pp. 171–6. DOI: 10.1053/apmr.2002.28025 (cit. on p. 3).

[8] Lawrence Ward. «Is Stochastic Resonance Just an Epiphenomenon?» In: 665 (May 2003). DOI: `10.1063/1.1584878` (cit. on pp. 3, 12).

[9] Kim Blackwell. «The Effect of White and Filtered Noise on Contrast Detection Thresholds». In: *Vision research* 38 (Jan. 1998), pp. 267–80. DOI: `10.1016/S0042-6989(97)00130-2` (cit. on pp. 3, 15).

[10] Yunfan Gong, Nestor Matthews, and Ning Qian. «Model for Stochastic-Resonance-Type Behavior in Sensory Perception». In: *Phys. Rev. E* 65 (3 Feb. 2002), p. 031904. DOI: `10.1103/PhysRevE.65.031904`. URL: `https://link.aps.org/doi/10.1103/PhysRevE.65.031904` (cit. on p. 3).

[11] Claudio Perez, Theodore Cohn, Leonel Medina, and José Donoso. «Coincidence-Enhanced Stochastic Resonance: Experimental Evidence Challenges the Psychophysical Theory Behind Stochastic Resonance». In: *Neuroscience letters* 424 (Sept. 2007), pp. 31–5. DOI: `10.1016/j.neulet.2007.07.014` (cit. on pp. 3, 15).

[12] D.M. Green and J.A. Swets. *Signal Detection Theory and Psychophysics*. Wiley, 1966. ISBN: 9780471324201. URL: `https://books.google.de/books?id=Ykt9AAAAMAAJ` (cit. on pp. 3, 4, 7).

[13] Jakob Tougaard. «Signal detection theory, detectability and stochastic resonance effects». In: *Biological cybernetics* 87 (Sept. 2002), pp. 79–90. DOI: `10.1007/s00422-002-0327-0` (cit. on p. 4).

[14] George A. Gescheider. *Psychophysics: The Fundamentals*. 3rd ed. Psychology Press, 1997. ISBN: 080582281X,9780805822816 (cit. on p. 4).

[15] Peter Newman (eds.) John Eatwell Murray Milgate. *Game Theory*. 1st ed. The New Palgrave. Palgrave Macmillan UK, 1989. ISBN: 978-0-333-49537-7,978-1-349-20181-5 (cit. on p. 5).

[16] Philip Higham and Michelle Arnold. «Beyond reliability and validity: The role of metacognition in psychological testing». In: Mar. 2007, pp. 139–162 (cit. on p. 6).

[17] David Pascucci, Giovanni Mancuso, Elisa Santandrea, Chiara Della Libera, Gijs Plomp, and Leonardo Chelazzi. «Laws of concatenated perception: Vision goes for novelty, decisions for perseverance». In: *PLOS Biology* 17 (Mar. 2019), pp. 1–42. DOI: `10.1371/journal.pbio.3000144` (cit. on p. 8).

[18] Luca Gammaitoni, Peter Hanggi, Peter Jung, and Fabio Marchesoni. «Stochastic Resonance». In: *Reviews of Modern Physics* 70 (Mar. 1999). DOI: `10.1103/RevModPhys.70.223` (cit. on p. 11).

[19] Jose Vilar, Gabriel Gomila, and J. Rubi. «Stochastic Resonance in Noisy Nondynamical Systems». In: *Physical Review Letters* 81 (July 1997), pp. 14–17. DOI: `10.1103/PhysRevLett.81.14` (cit. on p. 11).

[20] Lawrence Ward. «Psychophysics of Stochastic Resonance». In: *Fluctuation and Noise Letters* 04 (Jan. 2012), pp. L11–L21. DOI: 10.1142/S0219477504001616 (cit. on pp. 11, 12).

[21] Goran Soderlund and Sverker Sikström. «Positive effects of noise on cognitve performance: Explaining the Moderate Brain Arousal model». In: (June 2021) (cit. on p. 12).

[22] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. ISBN: 1441412697 (cit. on p. 16).

[23] Jonathan Peirce, Jeremy R. Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. «PsychoPy2: Experiments in Behavior Made Easy». In: *Behavior Research Methods* 51.1 (Feb. 2019), pp. 195–203. DOI: 10.3758/s13428-018-01193-y. URL: https://doi.org/10.3758/s13428-018-01193-y (cit. on p. 16).