Lecture 39 - The Theory of Signal Detection

THE THEORY OF SIGNAL DETECTION

Whenever a subject tries to detect a stimulus, his threshold criteria can vary depending on factors such as motivation, attention, fatigue, etc. These affect the response. The response can also vary depending on the amount of background noise in the neural system at that moment. The **theory of signal detection** tries to account for the influence of background neural noise and varying subjective criteria on the measured threshold.

Neural activation needed for detection

Whenever the visual system is stimulated, the brain receives an electrical signal. Schwartz refers to this as **neural activation**. For very weak stimuli, the neural activation is small, but for strong stimuli it is large. For the sake of explanation, let us assume that, in order for the subject to detect the stimulus (see it), the neural activation must exceed some fixed threshold criterion in the brain. Whenever the neural activation exceeds this threshold, the subject will think, “I see it.” For neural activation below this, the subject will think, “I can’t see it.” (Figure 1)

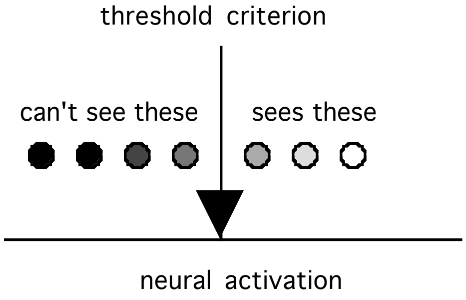

Figure 1. Whenever the retinal receptors receive an image, they send a signal to the brain that is proportional to the stimulus strength. Whenever the neural signal is strong enough, the person will perceive that he sees the object.

Background neural noise

Even when the stimulus is not present, neurons randomly fire, though the level of electrical activity will be very low. Nonetheless, it is always present and is called background **neural noise**. The strength of the neural noise is variable. Sometimes the noise creates a very small level of neural activation, sometimes more. At any moment, the value of the neural activation will vary randomly, but we can assume that the range of values will have a normal distribution about some mean value. Figure 2 shows a probability distribution for the values of the neural noise. It is a bell-shaped curve, which represents a normal, or Gaussian distribution. Even when no stimulus is present, the brain will receive neural activation with a probability distribution such as that shown in the graph.

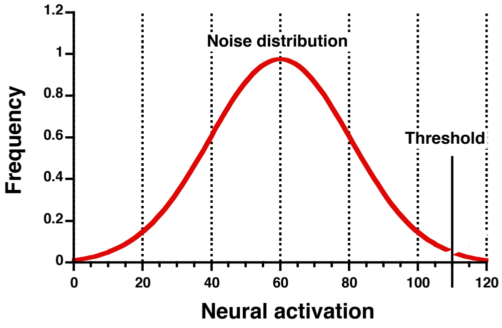

Figure 2. Distribution of neural activation due to noise alone

Figure 2 also indicates the level of neural activation for the threshold. Background neural noise usually produces a weak signal that is well below the threshold criterion for detection. However in some rare instances the background neural noise might be strong enough that it exceeds the threshold criterion. In that case the person would mistakenly think he is seeing the stimulus, but he is actually only receiving neural activation caused by neural noise. The probability of this happening is low. It is proportional to the *area under the curve to the right of the criterion line.* Of course, this can only occur if the threshold criterion is low enough that it falls within the range of the noise distribution.

Neural excitation when the stimulus is present

We will assume that, for one particular intensity setting, the stimulus alone always causes a fixed amount of neural activation whenever it is presented. For example, let us assume that the stimulus alone causes a neural signal with a strength of 100 arbitrary units. Since noise is always present, the total neural activation received when the *stimulus is on* will be due to the combined effects of the **stimulus+neural noise**.

Since the stimulus always has the same value, and the noise varies randomly, the sum of stimulus+noise will also vary. Since the variation is completely due to the changes in noise, the probability distribution of the stimulus+noise signal will be the same shape as the noise probability distribution (bell-shaped curve), but it will be shifted to the right. This is because excitation due to the stimulus is added to the noise. (See Figure 3 below or Schwartz Fig. 11-5.)

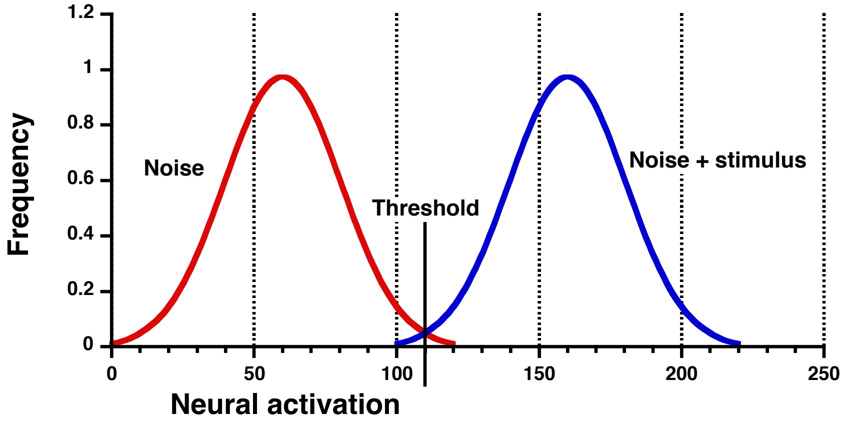

Figure 3. Two distributions; the left bell curve represents the probability distribution for noise alone. The right curve represents the distribution for the stimulus plus noise.

Note that the two curves show two different stimulus conditions—when the stimulus is turned off and when it is turned on. The subject will receive neural activation from either one or the other distribution. Similar plots are shown in Schwartz, but it is important to understand that the graph does not show two conditions present at the same time. Rather they are two possible presentations, either with the stimulus on or off, and both distributions are shown on the same plot for comparison. At any one time, the neural signal will come from either one or the other of these distributions.

Note where the threshold criterion is drawn, and consider the stimulus+noise curve only. Most of the time, when the stimulus is turned on, the neural activation from the stimulus plus noise will be above the criterion line, and the brain will respond, “I see it.”

The probability of seeing the stimulus, when it is on, is proportional to the area under the curve to the right of the criterion line. Assuming the total area is equal to 1.0, the probability of detection, in this case, would be approximately 98%.

Also note, however, that sometimes even when the stimulus is present, the noise is so low that the combined neural signal from the stimulus plus noise falls below the threshold criterion line. Whenever the neural activation is this low, the brain responds, “I see nothing.” This is indicated by the *area under the curve to the left of the criteria.* According to this curve, the probability of this happening is very low—about 2%.

Neural activation when no stimulus is present
When no stimulus is present, the probability distribution for receiving a certain level of neural activation (due to noise alone) is shown by the left bell shaped curve in Figure 3. Most of the time, the neural activation caused by the noise alone is below the threshold criterion, and the subject will correctly say that he sees nothing. The probability of correctly saying that the stimulus is not present, when it is not present, is *proportional to the area under the noise-only curve to the left of the criterion line*. In this example, there is a 98% probability that the person will correctly state that the stimulus is not present.

Also note that, occasionally the noise alone is sufficient to exceed the threshold criterion, and the patient will incorrectly think he is seeing the stimulus, but it is just background neural noise. The probability of this happening is proportional to the area under the noise-only curve to the right of the criterion line—in this case, about 2%.

Detectibility
A subject must decide whether or not he sees a stimulus based upon the neural excitation he receives. He will not know which distribution the neural excitation came from—whether the excitation came from the noise only, or if it came from the stimulus+noise. As long as it exceeds the threshold, the response is the same; he thinks, "I can see it."

On every trial in the experiment, the person tries to decide if he is looking at the noise-only distribution or the stimulus+noise distribution, but the only information available to him is the neural excitation value. If the stimulus intensity is small, the right curve will be shifted to the right of the noise-only curve only slightly and there will be much overlap of the two distributions. In that case it will be difficult to pick the stimulus out from the noise, and the subject will probably make many errors.

If the stimulus intensity is large, the stimulus+noise distribution will be pushed further to the right and it will be easier to distinguish between the distributions. The subject will probably make few mistakes.

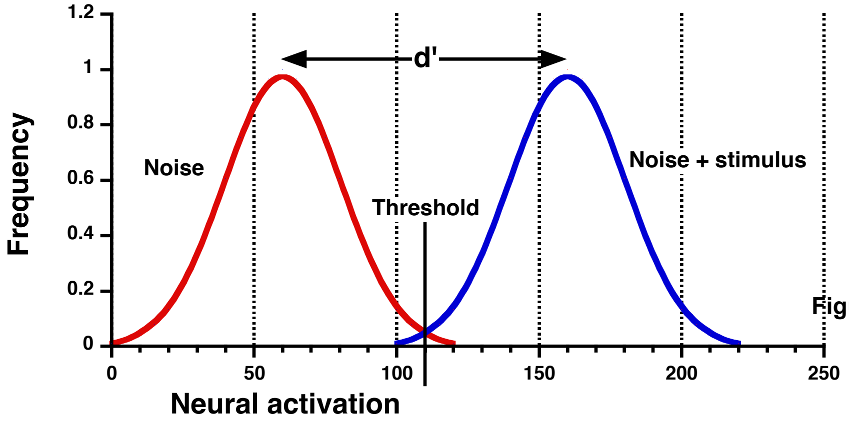

Figure 4. Detectibility (d')

The distance between the peaks (or mean values of the two probability distributions) indicates the strength of the stimulus intensity, and is called the **detectibility** (d' on the graph). When stimulus intensity is high and detectibility (d') is large, the stimulus will be easy to detect, and the subject will make few mistakes. When d' is small, the curves will have large areas of overlap. It will be harder for the subject to distinguish the stimulus+noise from noise alone, and he will make more mistakes. Refer to Schwartz Fig. 11-6.

EFFECT OF DIFFERENT CRITERIA
Schwartz Fig. 11-7 illustrates how a subject's response to a detection task will depend on the level of his criterion. In this example, d' is small (small stimulus intensity), so if the noise (N) and stimulus+noise (N+S) distribution were drawn on the same plot, they would have considerable overlap. It is important to remember, however, that during any particular trial, the subject will see a neural signal produced by EITHER

the N or N+S distribution, but not both. How would different criteria affect results according to the theory of signal detection?

Consider a hypothetical experiment. A subject sits in a perfectly dark room facing a panel and, after hearing a tone, he must report whether he saw the faint light or not. In order to keep him honest, the experimenter will mix in some null trials in with the presentations. That is, sometimes when the tone sounds, the light will not be turned on; other times it will be present. Every time he hears the tone, the subject must say, "I saw it" or "I didn't see it".

For every presentation, there are two possible responses, each of which could be right or wrong. Table 1 summarizes the four possible outcomes.

Table 1. Four possible outcomes in a detection experiment.

Stimulus / response	**Stimulus on**	**Stimulus off**
"Yes, I saw it."	Correct Hit (true positive)	Error (Type I) False alarm (false positive)
"No, I didn't see it."	Error (Type II) Miss (false negative)	Correct Correct rejection (true negative)

How would the subject respond for various criteria levels?

Lax criteria - noise only
Figure 11-7A shows how the subject would respond if the criterion were low or lax. The subject will say, "I see it." for low levels of neural excitation. On a visual fields test you would call such a person, "trigger happy."

Figure 11-7A, top, shows how the person would respond if no stimulus were present during a particular trial. In this case, the brain receives neural activation from the noise alone. Much of time excitation produced by N alone is above the criterion. He will say, "see it," but it's just noise. This response is called a **false alarm**, or **false positive**.

Q. What is the probability of getting a false alarm with this criterion?
A.

Sometimes excitation from N will be below the criterion. He responds that he sees nothing, which is correct—there is nothing but noise. This is called a **correct rejection** or **true negative**.

Q. What is the probability of getting a correct rejection with this criterion?
A.

Lax criterion - stimulus + noise
Figure 11-7A, bottom, shows how the person would respond if the stimulus were turned on during a presentation. The criterion is the same, but now the neural excitation is caused by the N+S distribution. Most of time, stimulus+noise will be above criterion. He says, "I see it", which is correct. This is called a **hit** or **true positive**.

Q. With this detectibility and this criterion, what is the probability of getting a hit?
A.

On a few occasions, the S+N will produce only weak neural excitation. The signal is delivered at a time when the background noise is very low. In spite of the fact that the stimulus is present, the subject responds, "I didn't see it." This is called a **miss** (**false negative**), since he failed to see the stimulus when it was present.

Q. With this detectibility and this criterion, what is the probability of getting a miss?
A.

From these examples, we can see that, *when the criterion is set low, you get many hits but also many false alarms.*

It is possible to influence a patient's criterion using penalties or rewards. For example, suppose you tell a subject, "Every time you see the dim light, I'll give you $100". This will encourage the patient to set his criterion low, and he will be quick to say "I see it." This will cause many false alarms, but he will also correctly say he sees it when it is on most of the time (many hits).

Strict criteria - noise only
Figure 11-7B, top, shows how the subject would respond if the criterion were strict. The subject will say, "I see it." only if the neural excitation is relatively high. Only on the few occasions when the noise exceeds the criterion will the person say he sees it.

Q. What is the probability of getting a false alarm (false positive) with this criterion?
A.

Most of the N distribution is below the criterion, so most of the time, when the stimulus is not present, the person will correctly say that he doesn't see it.

Q. What is the probability of getting a correct rejection with this criterion?
A.

Strict criteria - signal + noise
Figure 11-7B, bottom, shows how the subject would respond if the stimulus were turned on, while the criterion is strict. Part of the N+S distribution is above the criterion. This indicates the probability that the person would say that he sees it, when it is present (a hit).

Q. With this detectibility and this criterion, what is the probability of getting a hit?
A.

Most of the N+S distribution is below the criterion. This region represents the cases in which the subject will say he doesn't see it when it is actually present (a miss).

Q. With this detectibility and this criterion, what is the probability of getting a miss?
A.

With a strict criterion, you get few false alarms, but many misses (few hits). You can influence the subject to set his criterion high by penalties. For example, if you say, "I'll give you an electric shock If you say you see it when it's not really there." The subject will probably set his criterion very high and only respond if he is absolutely sure he sees it.

Using the same logic, you should analyze the probability of getting each of the four responses with the moderate criterion shown in Fig. 11-7C.

- False alarm (false positive)
- Correct rejection (true negative)
- Hit (true positive)
- Miss (false negative)

In both clinical and experimental tests, it is possible to modify a subject’s criteria by encouraging, coaxing or by offering rewards or penalties. Think of how you do this in clinical testing.