g

Chapter 3

Quantifying Coding Interactions

The previous chapter discussed in detail methods for estimating the information that spike trains of *single neurons* convey about stimuli. The current chapter describes an information theoretic approach to the study of high order correlations among small groups of neurons in the auditory system. Section 3.1 reviews previous studies based on the information theoretic approach to this problem, and section 3.2 discusses methodological issues involved in measuring and estimating synergy and redundancy. The application of these methods to neural coding in the auditory system are presented in the next chapter.

3.1 Previous work

Several investigators have used information theoretic measures for quantifying high order correlations in the activity of sets of neurons. These were mostly performed in the visual system, but also in high brain areas [Gawne and Richmond, 1993, Meister, 1996, Warland *et al.*, 1997, Rolls *et al.*, 1997, Dan *et al.*, 1998, Gat and Tishby, 1999, Brenner *et al.*, 2000, Nirenberg *et al.*, 2001, Reich *et al.*, 2001].

Gawne and Richmond (1993) defined information theoretic measures to quantify correlations in pairs of inferior temporal visual neurons. They found both independent, redundant and synergistic pairs, and concluded that IT

cortical pairs were not organized in clusters of similar-properties neurons, but were heterogeneous and more independent than redundant.

Rolls, Treves and Tovee (1997), measured the information conveyed by set of neurons of varying sizes, observing an almost linear growth of information as a function of number of neurons, as expected under a distributed coding scheme.

Warland and colleagues [Warland *et al.*, 1997] compared information extracted using the method of [Bialek *et al.*, 1991] using two retinal ganglion cells with the information conveyed when neurons were considered independently. Their main finding was that the traditional classification into ON/OFF types determines the synergistic/independence nature of neuronal correlations.

Gat and Tishby (1999) measured synergy and redundancy in pre-motor cortical neurons, in the context of Synfire chains. Their major finding was high synergy values that were observed in cortical but not in basal ganglia neurons.

Brenner and colleagues [Brenner *et al.*, 2000] measured information carried by compound events in spike trains of single neurons, and measured synergy and redundancy in their code. They showed how pairs of spikes can be synergistic or redundant in a way that depends on their time interval.

Recently, Nirenberg and colleagues [Nirenberg *et al.*, 2001] measured information levels in pairs of isolated ganglion neurons from the mouse retina. The main goal of their work was to estimate the relative weight of information transmitted by temporal correlations activity in coding of natural images. They found that more than 90 percent of the information could be extracted from the neurons when their temporal correlations were ignored. One should note however that their conclusion, "retinal ganglion cells act largely as independent encoders" contradicts earlier studies [Meister *et al.*, 1995, Meister, 1996]. Regarding the issues discussed in the current chapter, it worth emphasizing that Nirenberg and colleagues define *excess correlation function* (ECF) as a measure of the information lost when estimating information under stimulus conditioned independence $P(r_1, r_2|s) = P(r_1|s)P(r_2|s)$, and find that the ECF is small in their preparation.

Finally, Reich and colleagues [Reich *et al.*, 2001] measured redundancy levels in clusters of up to six simultaneously recorded neurons in primary visual cortex. Their main finding is that responses were almost independent under a *labeled lines* coding scheme, that is, when keeping track which neuron fired which spikes, but were redundant if responses were summed over clusters. This suggests that summing neuronal responses in a naive manner (for example in order to reduce variability) discards important information about the stimuli.

The common objective of this series of studies is to identify the way neurons interact to convey sensory information. This goal yields further refined questions: Are nearby neurons synergistic, independent or redundant? How should these qualitatively different behaviors be measured and quantified? Do these properties change along the ascending sensory pathways? Are they stimulus dependent? Perhaps most importantly, can such interactions be used by a readout mechanism and account for behavioral changes?

The current chapter aims at answering some of these questions. It has two main goals. First, from the methodological point of view, it will systematically define information theoretic measures of correlations between groups of cells, and discuss how these should be reliably estimated and compared. Then, from a scientific point of view, we will use these methods for the analysis of auditory neurons responses, to discover principles governing interactions between neurons in the auditory pathway. The unique nature of our data, namely neural responses from three brain regions to the same set of stimuli, allows us to compare the way typical correlations change along the auditory pathway.

3.2 Measures of synergy and redundancy

3.2.1 Preliminaries: synergy and redundancy in pairs

Let X_1 and X_2 be a pair of neurons conveying information about a stimulus S. Their synergy-redundancy measure is commonly defined as the difference between the amount of information that can be obtained when the two neurons area are considered jointly and the information obtained when they

are considered individually:

$$SR_{pairs}(X_1, X_2, S) = I(X_1, X_2; S) - [I(X_1; S) + I(X_2; S)] \tag{3.1}$$

Intuitively, SR_{pairs} measures the amount of information on the stimulus S gained by observing the joint distribution of both X_1 and X_2, as compared to observing each of the two cells independently. In the extreme case where $X_1 = X_2$, the two cells are completely redundant and provide the same information about the stimulus, yielding $SR_{pairs} = I(X_1, X_2; S) - I(X_1; S) - I(X_2; S) = -I(X_1; S)$, which is always non-positive. On the other hand, positive SR_{pairs} values testify to synergistic interaction between X_1 and X_2. For example, let X_1 and X_2 be fair Bernoulli variables and S their sum modulu 2 ,$S = X_1 \oplus X_2$. In this case any isolated variable X conveys zero information about S while knowing their joint value provides all the information about S.

Although the SR_{pairs} measure is defined in an asymmetric way, it obeys the following property

Lemma 3.2.1: *The SR_{pairs} is symmetric with respect to all three variables*

$$\begin{aligned} SR_{pairs} &= I(X, Y; Z) - I(X; Z) - I(Y; Z) \\ &= I(Z, Y; X) - I(Z; X) - I(Y; X) \\ &= I(X, Z; Y) - I(X; Y) - I(Z; Y) \end{aligned} \tag{3.2}$$

Proof: First, use the chain rule for the mutual information (section **??**) $I(X, Y; Z) = I(X; Z) + I(Y; Z|X)$ and write

$$I(X, Y; Z) - I(X; Z) - I(Y; Z) = I(Y; Z|X) - I(Y; Z) \ . \tag{3.3}$$

Secondly, from the symmetry in the MI definition $I(X, Y; Z) = I(Y, X; Z)$, we obtain $I(X, Y; Z) = I(X; Z|Y) - I(X; Z)$. To complete the triple symmetry, we write

$$\begin{aligned} I(X, Y; Z) &- I(X; Z) - I(Y; Z) \\ &= H(X, Y) - H(X, Y|Z) - H(X) + H(X|Z) \\ &\quad -H(Y) + H(Y|Z) \\ &= I(X; Y|Z) - I(X; Y) \end{aligned} \tag{3.4}$$

□

As a conclusion from this derivation, the measure SR_{pairs} can be written as the difference of two MI terms, a stimulus conditioned term and an unconditioned one

$$\begin{aligned} SR_{pairs} &= I(X_1, X_2; S) - [I(X_1; S) + I(X_2; S)] \\ &= I(X_1; X_2|S) - I(X_1; X_2) \end{aligned} \tag{3.5}$$

The first term in the last expression is commonly referred to in the literature as *noise correlations* [Gawne and Richmond, 1993, Panzeri *et al.*, 1999], measuring the level of dependence given the stimuli, while the second is known as *signal correlations*, measuring the dependencies induced by the different stimuli. This representation makes it easier to see how SR_{pairs} can assume both positive and negative values. Since the positive values of SR_{pairs} are commonly interpreted as synergy, and the negative values as redundancy, we naturally treat $I(X_1, X_2|S)$ as measuring the synergistic interaction of the pair, while $I(X_1; X_2)$ measures their redundancy interactions. Separating SR_{pairs} into the difference of the two terms allows us to measure separately the two different effects.

The case of synergistic coding, $SR_{pairs} > 0$, means that specifying the stimulus increases the MI between two units, and makes a precise definition of the idea of "stimulus dependent correlations" [Abeles *et al.*, 1993, Vaadia *et al.*, 1995, Hopfield, 1995, Meister *et al.*, 1995, Meister, 1996, Singer and Gray, 1995, Brenner *et al.*, 2000]. However, care should be taken when interpreting the $I(X_1, X_2; S)$ as standing for the everyday meaning of the word synergy. For example, consider two independent binary variables whose activity is contaminated with zero mean correlated noise ξ that is independent of the stimulus, e.g. with a joint distribution

$$P(X, Y|S = s) = \begin{pmatrix} 0.25 + \xi & 0.25 - \xi \\ 0.25 - \xi & 0.25 + \xi \end{pmatrix}. \tag{3.6}$$

In this case the synergy term is strictly positive $I(X; Y|S) > 0$ but the correlated noise is not actually used for coding (since the noise is stimulus independent).

3.2.2 Estimation considerations

In chapter 2 we derived the bias of the mutual information estimator, showing that it is roughly proportional to the ratio between the number of free parameters estimated and the number of samples. We now use this result to derive the bias of the synergy and redundancy terms, and show that these biases are considerably different. To see this, denote the total number of samples n, the number of samples per stimulus n_s, the number of stimuli $|S|$, the number of possible responses $|R|$, and stimulus probabilities $P(S = s) = n_s/n$. Since the possible responses of the two neurons reside in a two dimensional matrix of size $|R| \times |R|$, the stimulus conditioned joint probability matrix has n_s samples in $|R|^2$ bins for each stimulus, yielding a bias of $\frac{|R|^2}{2n_s \log(2)}$ per stimulus. Thus the total bias of the synergy term equals

$$E[I(X;Y|S)] = \sum_s P(S = s) \frac{|R|^2}{2n_s \log(2)} = \frac{|S||R|^2}{2n \log(2)} \tag{3.7}$$

On the other hand, the bias of the redundancy term is S times smaller

$$E[I(X;Y)] = \frac{R^2}{2n \log(2)} = \frac{1}{|S|} E[I(X;Y|S)] \tag{3.8}$$

As a consequence, the bias of the synergy-redundancy estimator is again positive, and equals

$$E[I(X;Y|S) - I(X;Y)] \quad = \quad \frac{(|S| - 1)|R|^2}{2n \log(2)} \tag{3.9}$$

If the mutual information estimators are not properly corrected for the bias, independent neurons will appear as synergistic ones. This effect is not specific to estimators that use a discretizing procedures into bins, but stems from the fact that the synergy term requires estimating a finer effect, with mutual information terms with a smaller relative number of samples than the redundancy term.

3.2.3 Extensions to group redundancy measures

We now turn to extend the measures of synergy and redundancy beyond neuronal pairs. In the multivariate case, several different definitions of synergy and redundancy naturally arise, and each measures a different effect.

Redundancies of N-tuples given singles

First, as in the case of pairs, one may be interested in the difference between information levels conveyed by the joined distribution of N variables (neurons) compared to that provided by N single independent ones. This is defined by

$$SR_{N|1} = I(X_1, ..., X_N; S) - \sum_{i=1}^{N} I(X_i; S) \tag{3.10}$$

As with SR_{pairs}, this synergy-redundancy measure may be rewritten as the difference between two multi-information terms (see definition **??**)

$$\begin{aligned} SR_{N|1} &= I(X_1, ..., X_N; S) - \sum_{i=1}^{N} I(X_i; S) = \\ &= H(X_1, ..., X_N) - H(X_1, ..., X_N|S) \\ &\quad - \sum_{i=1}^{N} H(X_i) - H(X_i|S) = \\ &= I(X_1; ...; X_N|S) - I(X_1; ...; X_N) \end{aligned} \tag{3.11}$$

The index $SR_{N|1}$ can thus be separated into two terms, as with SR_{pairs}. The first is always non-negative, measures *within-stimulus correlations* (noise correlations) and is again termed here *synergy*. The second is always non-positive, measures *between stimulus correlations*, (signal correlations), and again quantifies the *redundancy*.

A major difficulty in estimating $SR_{N|1}$ is that it requires estimating the joint distribution of $N + 1$ variables: the stimulus and the activity of N neurons. This often becomes prohibitive even for the small N used in electrophysiological experiments, and is discussed in section 3.3.1.

Redundancies of N-tuples given N–1-tuples

In contrast to the two-variables case, the multi-variable case enables other measures of redundancies. For example, one may be interested in the difference between information conveyed by all N variables and that conveyed by pairs, triplets, or even all the $N-1$ subsets of variables. A case of particular interest is the residual information obtained from the joint distribution of

all N variables, as compared to that obtained from any subset of $N-1$ variables. As with inclusion-exclusion calculations, the SR measure for this case is defined as

Definition 3.2.2 : N**-tuples given** N–1**-tuples synergy-redundancy**

$$\begin{aligned} SR_{N|N-1} &= I(X^{(N)};S) \\ &\quad - \sum_{X^{(N-1)}\in I^{N-1}} I(X^{(N-1)};S) \\ &\quad + \sum_{X^{(N-2)}\in I^{N-2}} I(X^{(N-2)};S) \\ &\quad \vdots \\ &\quad +(-1)^{N-1}\sum_{\{X_i\}} I(X_i;S) \end{aligned} \tag{3.12}$$

where I^k is the set of all subsets of size k out of the N variables.

For example, for $N=3$ we have

$$\begin{aligned} SR_{3|2} &= I(X_1,X_2,X_3;S) \\ &\quad -I(X_1,X_2;S) - I(X_1,X_3;S) - I(X_2,X_3;S) - \\ &\quad +I(X_1;S) + I(X_2;S) + I(X_3;S) \end{aligned} \tag{3.13}$$

This definition contains a total of 2^N-1 terms. For $N=2$ it coincides with $SR_{pairs} = SR_{2|1}$.

As with the previous redundancy measures, it can again be rewritten as a difference between conditional and unconditional multi-information terms

$$\begin{aligned} SR_{N|N-1} &= I(X^N|S) - \sum_{X^{N-1}\in I^{N-1}} I(X^{N-1}|S) + ... \\ &\quad -I(X^N) + \sum_{X^{N-1}\in I^{N-1}} I(X^{N-1}) - ... \end{aligned} \tag{3.14}$$

where in this context $I(X^N) = I(X_1;...;X_N)$, and I^k is defined as in 3.12.

Calculating $SR_{N|N-1}$ is difficult in practice since it requires estimating $O(2^N)$ information terms. This difficulty is compounded by the difficulty encountered in estimating of $SR_{N|1}$ in which joint probabilities of exponential sizes have to be estimated. These issues are discussed in the following section.

3.3 Redundancy measurements in practice

Applying the measures described above to finite data, and comparing redundancies across different brain stations involves several delicate points. We discuss three of them here: using stimulus conditioned independence approximation, the effect of baseline single-unit information level and redundancy bias due to information ceiling effects.

3.3.1 Conditional independence approximations

The multi-variable redundancy measures $SR_{N|1}$ and $SR_{N|N-1}$ defined in the previous section are based on a joint distribution of $N+1$ variables $X_1, \ldots, X_N$ and S. Unfortunately, given the typical recordings in electrophysiological experiments, the sample sizes are rarely sufficient for reliable estimation of these joint distributions and consequently the MI values, even for moderate N values. One approach is to approximate this joint distribution with other distributions, using some predefined assumptions on conditional independence between variables. This approach has been advocated in the recent years in the literature of graphical models (Bayes nets), where predefined (conditional) independence structures are conveniently represented in the form of graphs [Pearl, 1988, Jordan, 1998, Shafer and Pearl, 1990].

One such approximating assumption is that neural activity of different neurons is conditionally independent given the stimulus

$$\begin{aligned} p(x_1, ..., x_N|S=s) &= \prod_{i=1}^{N} p(x_i|S=s) \qquad (3.15) \\ p(x_1, ..., x_N) &= \sum_s p(s) \left(\prod_{i=1}^{N} p(x_i|S=s) \right) . \end{aligned}$$

Figure 3.1 shows a graphical model illustration of this dependence structure.

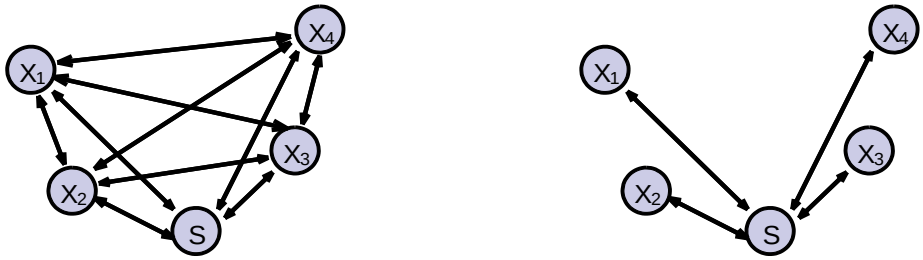

Figure 3.1: An illustration of dependence structure for two cases. **Left:** Four neurons and the stimulus, without any independence approximation. **Right:** Four neurons under the stimulus conditioned independence assumption.

Figure 3.2 depicts the joint distribution of the neural responses (number of spikes) of two IC neurons calculated under the conditional independence approximation. Each stimulus conditioned distribution was calculated under independence approximation

$$p(r_1, r_2) = \sum_s p(r_1, r_2|s)p(s) = \sum_s p(r_1|s)p(r_2|s)p(s). \tag{3.16}$$

Estimating the full joint distribution $p(R_1, R_2)$ thus requires estimating S distributions of size $|R_1|$ and S distributions of size $|R_2|$, instead of S two dimensional distributions of size $|R_1| \times |R_2|$.

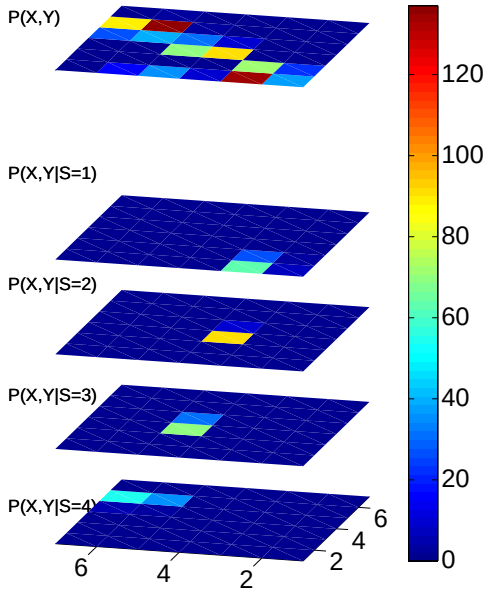

Figure 3.2: Stimulus conditioned joint distribution of spike counts for two IC cells. Each of the four lower pallets correspond to the distribution of spike counts for a different stimulus, under stimulus conditioned independence. The upper panel is the sum of conditional distributions over 15 stimuli. The color code denotes the number of spike counts events.

With stimulus conditioned independence, calculating the joint distribution used in the information terms of $SR_{N|1}$ is considerably simplified, as for each stimulus s we only have to estimate the marginal distributions $p(X_i|S=s)$ instead of the full distribution $p(X_1,...,X_N|S=s)$. It thus allows us to estimate an exponentially smaller number of parameters, which in the case of neuro-physiological data makes it feasible to estimate the MI in these approximated distributions. This approximation thus makes it possible to investigate redundancy among considerably larger groups of neurons than the 2-3 neuron groups often analyzed in the literature. In terms of the bias variance tradeoff discussed in Chapter 2, using this approximation decreases the variance, since it decreases the number of degrees of freedom, but increases the bias, as the estimated MI deviate from the ones calculated with a non approximated joint distribution.

Interestingly, under stimulus conditioned independence, the synergy term

$SR_{N|1}$ vanishes, thus limiting neural interactions to the redundant (negative) regime. The measure we will use often in the current work is therefore

$$Redundancy_{N|1} = -I(X_1; \ldots; X_N) \tag{3.17}$$

We will also refer at times to the quantity $I(X_1; \ldots; X_N)$ as redundancy, but the difference should be clear from the context.

One situation in which the assumption of stimulus conditioned independence is crucial is in the analysis of non simultaneously recorded data. In this case, if we want to obtain any estimate of informational coupling between neurons, it is necessary to couple the given marginal distributions of the single neurons conditioned on the stimuli. Stimulus-conditioned independence is then the most natural way (in some sense, the maximum-entropy way) of performing this coupling.

How reasonable is the conditional-independence approximation? Naturally, its validity depends on the data at hand. In theory the approximation is expected to be good when neuronal activity is mostly determined by the presented stimulus and to a lesser extent by interactions with nearby neurons. One example is the high input regime of cortical neurons receiving thousands of inputs, where a single input has only a limited influence on the activity of the target cell. The experimental evidence in this regard is however mixed (see e.g.[Gochin *et al.*, 1994]).

When simultaneously recorded neurons are available, it is possible to test empirically the quality of the approximation for small sets of neurons (e.g. pairs and triplets), and quantify the deviation from independence. This was done e.g. in [Nirenberg *et al.*, 2001]. We performed these tests for a limited set of simultaneously recorded neurons in AI and MGB and found that the approximation causes only a small deviation from the true distribution. These results are described in detail in section 4.1 (see Fig. s 4.7 and 4.8).

3.3.2 The effect of single-unit information on redundancy

A major goal of the current work is to compare coding characteristics across different brain regions in a sensory pathway. Clearly, we must ensure that any observed difference in redundancies between regions is not an artifact

due to differences in other factors. In particular, the redundancy between neurons tends to grow when single-unit information about the stimulus grows. The differences in MI levels in different auditory stations therefore require normalizing the redundancies to a unified scale.

To demonstrate this effect, consider the following toy model (Fig 3.3), consisting of a Markov chain of three binary variables $X \rightarrow S \rightarrow Y$, where each of the pairs X, S and S, Y is a binary symmetric noisy channel with noise level $\xi \in [0, 1]$. Formally we have, $Pr(X = 1) = \frac{1}{2}$ and $Pr(X = S) = Pr(S = Y) = 1 - \xi$. Since X, S and Y form a Markov chain, X and Y are conditionally independent given the stimulus S.

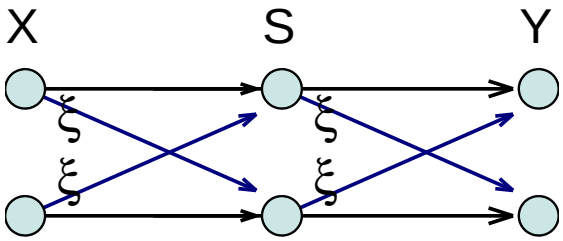

Figure 3.3: Two concatenated symmetric noisy binary channels, each with noise level ξ.

The intuitive interpretation for the notion of redundancy between X and Y, suggests that in this model redundancy should not depend on the noise level ξ. However, the noise level here effects both the single-unit information and the redundancy, as can be shown analytically, by calculating the single-unit information $I(X;S)$,$I(Y;S)$ and redundancy magnitude $I(X;Y)$ as a function of ξ

$$I(X;S) = I(S:Y) = 1 - H[\xi] \tag{3.18}$$

where $H[p]$ is the entropy of the binary distribution $(p, 1-p)$ $H[p] = -p\log_2 p - (1-p)\log_2(1-p)$. In addition, the pair (X, Y) is also a symmetric binary channel, this time with a noise level of $2\xi(1-\xi)$. The redundancy magnitude, which is the information of the concatenated channel, is therefor

$$I(X;Y) = 1 - H[2\xi(1-\xi)] \tag{3.19}$$

Figure 3.4A plots the single-unit information $I(X;S)$ and the redundancy magnitude $I(X;Y)$ as a function of the noise level ξ. It shows that although the redundancy changes considerably with the noise level it is roughly consistent with the change in the single-unit information. This suggests that the single-unit information provides a natural scale for the problem and can be used to normalize the redundancy measure. The relation between MI and redundancy as a function of the noise level is presented as a scatter plot in Fig. 3.4B.

These data suggest that the *normalized redundancy* $I(X;Y)/I(X;S)$ (or, which is equivalent in this case $2I(X;Y)/[I(X;S)+I(S;Y)]$) provides a reasonable normalization scheme for this simple example, because it preserves redundancy values $I(X;Y)$ over a range of single-unit information values $I(X;S)$. This normalization procedure was indeed used in other studies of synergy and redundancy [Brenner *et al.*, 2000, Reich *et al.*, 2001].

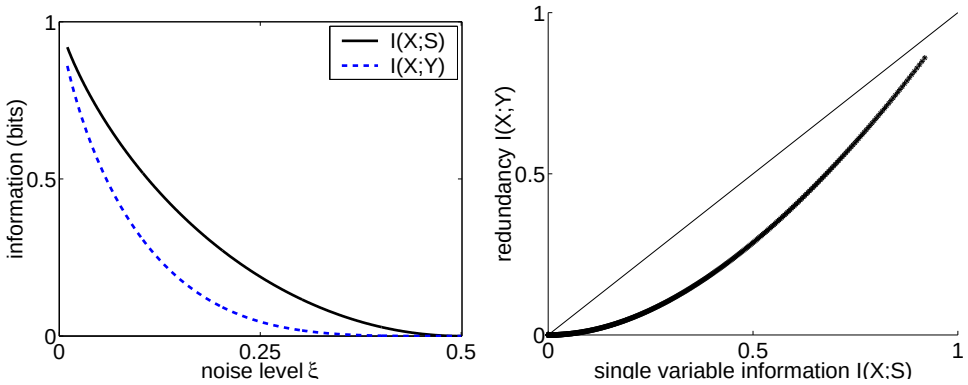

Figure 3.4: **left:** Mutual information and redundancy as a function of the noise level ξ. **right:** Scatter plot of MI vs. (non normalized) redundancy.

Consider now a second setting, where the Markov chain $X \rightarrow S \rightarrow Y$ consists of three variables, each having an alphabet of 4 letters. We compare two models, described in figure 3.5, where the intuitive notion of redundancy asserts that the model A (orthogonal noise in $p(x,s)$ and $p(y,s)$ is less redundant than the model B. For example, when the noise is maximal $\xi = \frac{1}{2}$, X and Y provide the same information on S in model B, since they serve to discriminate between the upper two alternatives, but different

information on S in model A.

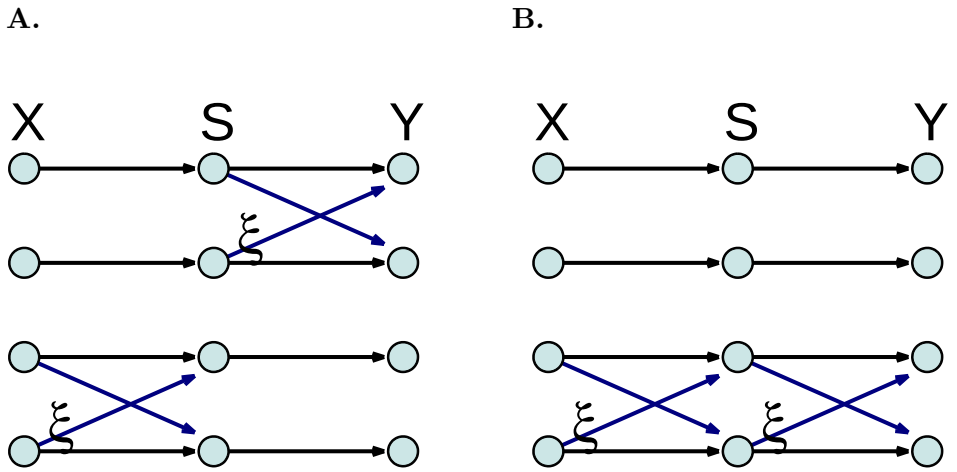

Figure 3.5: Two toy example for 4-letter channels, with different structure. Each channel consists of a concatenation of two channels, in which two of the values are transmitted with no noise, while the other two are corrupted with noise at level ξ. On the left channel, X and Y carry noiseless information about different pairs of values of S, while in the right panel both X and Y convey full information about the same pair of values of S.

Once again mutual information can be calculated analytically in these two models. For uniform $p(x)$ we have

$$\begin{aligned}
&\textbf{model A}: && && (3.20)\\
&I(X;S) &&= H(S) - H(S|X) = 2 - \frac{1}{4}\left(0 + 0 + H[\xi] + H[\xi]\right)\\
&&&= 2 - \frac{1}{2}H[\xi]\\
&I(X;Y) &&= H(Y) - H(Y|X) = 2 - \frac{1}{4}(H[\xi]+H[\xi]+H[\xi]+H[\xi])\\
&&&= 2 - H[\xi]\\
&\textbf{model B}:\\
&I(X;S) &&= H(S) - H(S|X) = 2 - \frac{1}{2}H[\xi]\\
&I(X;Y) &&= H(Y) - H(Y|X) = 2 - \frac{1}{2}H[2\xi(1-\xi)]
\end{aligned}$$

We now compare redundancies in these two models under different noise levels. Specifically, we wish to avoid a situation where one model seems

more redundant just because its single-unit information is higher. Figure 3.6 plots the single-unit information and redundancy in the two models as a function of noise level ξ. In both models the single-unit information and unnormalized redundancies are correlated (upper panels). More importantly, The normalized redundancy of model A is usually lower than that of model B (lower panels).

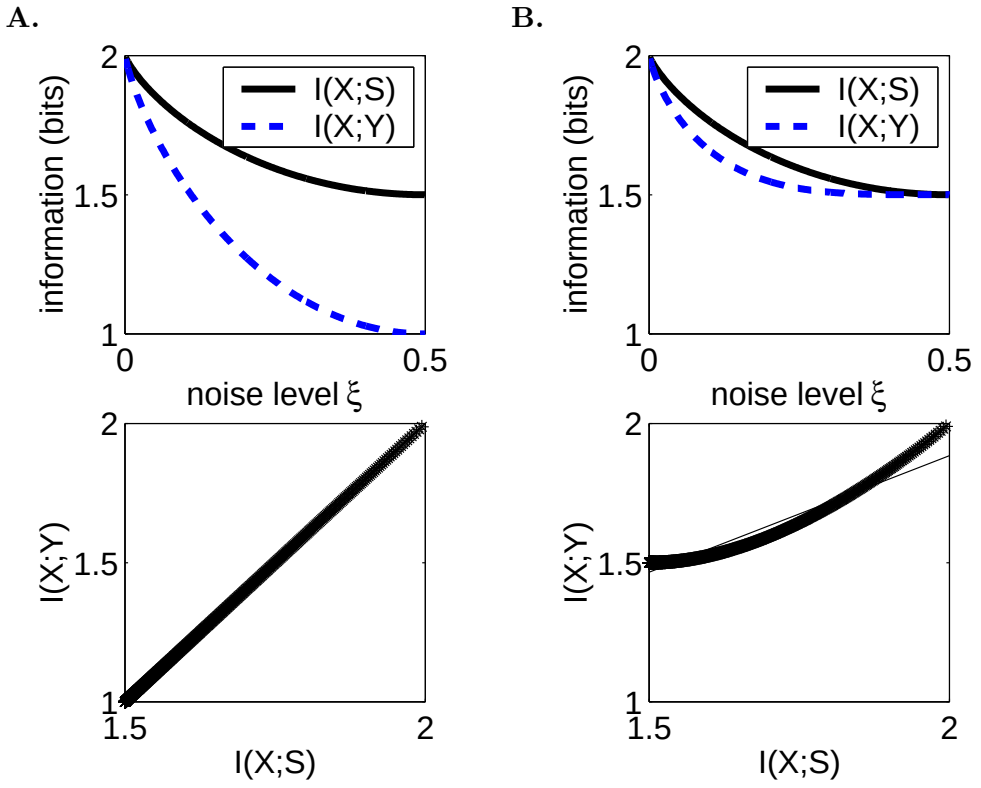

Figure 3.6: **Top:** Single information and redundancy in two synthetic models as a function of noise level. **Bottom:** Redundancy as a function of single-unit information, for the two models.

Figure 3.7 compares the normalized and non-normalized redundancies in the two models. For each of the two measures there is a regime of noise levels $\xi \in [0, \xi^{max}]$, for which any type-A model appears less redundant than any type-B model, regardless of ξ chosen independently for each model. This regime is almost three times larger for the normalized redundancy measure

than the unnormalized one.

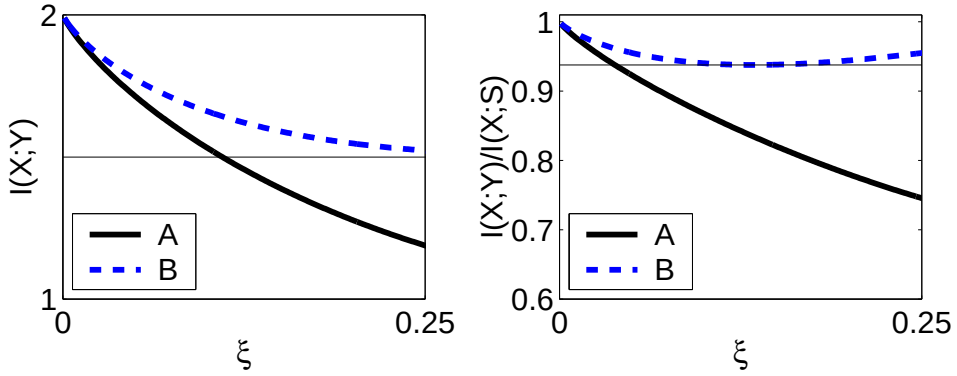

Figure 3.7: Comparing normalized and unnormalized redundancy measures in model A and B. The regime of ξ for which any type-A model appears less redundant than any type-B model is depicted with a thin line. This regime is about three times larger for the normalized redundancy measure. **A.** Non-normalized redundancy **B.** Normalized redundancy.

Clearly, these results do not guarantee that under more complex scenarios the redundancy will grow near-linearly with single neuron information. Naturally, the picture becomes even more complex for multi-neuron redundancies. The above results show however that normalizing by single neuron information captures the intuitive notion of redundancy better.

What about normalization of redundancy among larger groups of variables? Earlier we defined the redundancy in this case to be $-I(X_1;\ldots;X_N)$. This information is always bounded from above by the information between all variables and the stimulus

$$I(X_1;\ldots;X_N) \leq I(X_1;\ldots;X_N;S) \ , \tag{3.21}$$

and under stimulus conditioned independence this bound equals the sum of single unit information terms

$$\begin{aligned} & I(X_1;\ldots;X_N;S) = \\ &= \sum_{x_1,\ldots,x_N,s} p(x_1,\ldots,x_N,s) \log\left(\frac{p(x_1,\ldots,x_N,s)}{p(x_1)\cdot\ldots\cdot p(x_N)\cdot p(s)}\right) \\ &= \sum_{x_1,\ldots,x_N,s} p(x_1|s)\cdot\ldots\cdot p(x_N|s)p(s) \log\left(\frac{p(x_1|s)}{p(x_1)}\cdot\ldots\cdot\frac{p(x_N|s)}{p(x_1)}\right) \end{aligned} \tag{3.22}$$

$$
\begin{aligned}
&= \sum_{x_1,\ldots,x_N,s} p(x_1|s)\cdot\ldots\cdot p(x_N|s)p(s)\sum_{i=1}^{n}\log\left(\frac{p(x_i|s)}{p(x_i)}\right) \\
&= \sum_{i=1}^{n}\sum_{x_i,s} p(x_i|s)p(s)\log\left(\frac{p(x_i|s)}{p(x_i)}\right) \\
&= \sum_{i=1}^{n} I(X_i;S)
\end{aligned}
$$

Interestingly, tighter bounds such as the average or minimal single unit information do not hold. To see this consider the following example. Let $X_1,\ldots,X_N$ be symmetric binary variables, whose values are determined by the value of a binary symmetric S.

$$
(X_1,\ldots,X_N) = \begin{cases} (1,\ldots,1) & \text{when } S=1, \text{with } p=\frac{1}{2} \\ (0,\ldots,0) & \text{when } S=0, \text{with } p=\frac{1}{2} \end{cases} \tag{3.23}
$$

In this case, each single-unit information equals one $I(X_i;S) = 1$ bit $\forall i$. Thus both the average and the minimum of the single unit information terms are also 1 bit. However, the multi-information term grows with N, and equals

$$
\begin{aligned}
I(X_1;\ldots;X_N) &= \sum_{x_1,\ldots,x_N} p(x_1,\ldots,x_N)\log\left(\frac{p(x_1,\ldots,x_N)}{p(x_1)\cdot\ldots\cdot p(x_N)}\right) \qquad (3.24)\\
&= \sum_{s=0}^{1} p(\vec{X}=s)\log\left(\frac{p(\vec{X}=s)}{p(x_1=s)\cdot\ldots\cdot p(x_N=s)}\right) \\
&= \sum_{s=0}^{1}\frac{1}{2}\log\left(\frac{\frac{1}{2}}{\frac{1}{2}^N}\right) \\
&= N-1
\end{aligned}
$$

This example suggests that in the multivariate case redundancy should be normalized by the sum of the single unit information terms, which for a large N is a tight bound on redundancy. Therefore in this work we use the normalized multivariate redundancy measure

$$
\textit{Normalized Redundancy}_{\ N|1} = \frac{-I(X_1;\ldots;X_N)}{\sum_{i=1}^{n} I(X_i;S)}\ . \tag{3.25}
$$

Note that this measure of redundancy is equal up to a constant to the normalization examples discussed earlier for two variables (figures 3.3, 3.5).

3.3.3 Bias in redundancies estimation due to information ceiling effects

To illustrate the subject of the current section, we start with a simple example. Let $X_1, ..., X_N$ be the activities of N neurons in response to the stimuli, and let us assume they were measured independently, and that all of them have the same response distribution with the stimulus $P(X_i, S)$.

According to the chain rule for mutual information, the information conveyed by such independent variables about a stimulus $I(X_1, ..., X_N; S)$ increases according to $I(X_1, ..., X_N; S) = \sum_{i=1}^{N} I(X_i; S|X_1...X_{i-1})$. On the other hand however, this information is bounded by the stimulus entropy $H(S)$, which for any finite number of stimuli must be finite as well. This shows that the information accumulate sub-additively

$$\begin{aligned} I(X_1, ..., X_N; S) &= \sum_{i=1}^{N} I(X_i; S|X_1...X_{i-1}) \\ &< \sum_{i=1}^{N} I(X_i; S) = N\ I(X_1; S). \end{aligned} \tag{3.26}$$

The reason for this sub-additivity is that given the neurons $X_1, ..., X_{k-1}$ the information that the i^{th} neuron conveys about S is smaller than the unconditional information $I(X_k; S|X_1...X_{k-1}) < I(X_k; S)$. Simply stated, after some neurons convey information about the stimulus, there isn't much left to tell, and the remaining neurons can only convey lower levels of information. This effect is important when aiming to quantify redundancies, since it creates an "artificial" source of redundancy between neurons, stemming from the experimental conditions (namely the bounded entropy of the stimulus set).

How can this effect be quantified? A first order model of this phenomenon was suggested by [Gawne and Richmond, 1993, Rolls *et al.*, 1997]. Consider an abstract information "space" of size I^{max}, and N variables each conveying $I(X_i; S)$ bits of information, thus covering a fraction $I(X_i; S)/I^{max}$ of the space. The mean information conveyed by a single neuron is therefore $I_1 = \frac{1}{N} \sum_i I(X_i; S)/I^{max}$. If the information conveyed by each variable randomly covers some fraction of this space, the spurious overlaps between the any pairs of variables will be I_1^2. The expected fraction of information

covered by N variables is then

$$\begin{aligned} I(N) &= I^{max}\left[\binom{N}{1} I_1 - \binom{N}{2} I_1^2 + \binom{N}{3} I_1^3 - \ldots\right] \\ &= I^{max}\left(1-(1-I_1)^N\right) \ . \end{aligned} \tag{3.27}$$

In the limit of infinite number of neurons $N \to \infty$, $I(N)$ reaches the upper bound I^{max}.

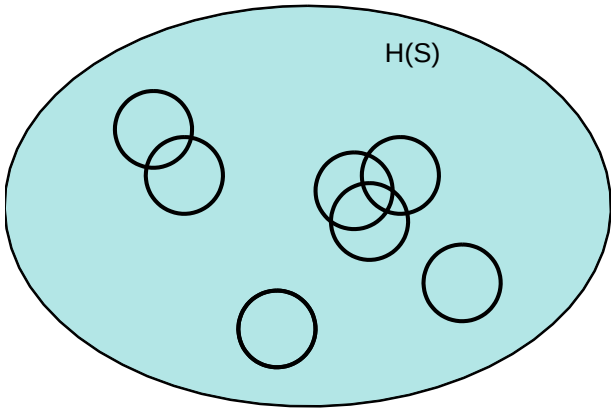

Figure 3.8: An illustration of the information space mode. The ellipsoid stands for the entropy of the stimulus $H(S)$, which is an upper bound on the maximal information an ensemble of neurons can convey about the stimulus. Each circle corresponds to a neuron, and its area corresponds to the mutual information $I(X_i;S)$. In this type of diagrams, neurons may in principle cover areas outside the ellipse, but these are not plotted here.

This model is loosely based on Venn type diagrams that correspond to entropies and mutual information through what is known as *I-measures* [Yeung, 1997]. For a more extensive discussion about this type of diagrams, see e.g. [Cover and Thomas, 1991] p.20 and [Csiszar and J.Korner, 1997] p.50.

The arguments using the information plane have intuitive appeal, but the theoretical justification behind them is not formally established. However, Samengo and Treves [Samengo and Treves, 2000, Samengo, 2001] compared the prediction of the information plane approach with actual information

curves for several synthetic examples and showed that the resulting equations reasonably approximate the observed information growth as a function of N. We therefore adopted this model as the null hypothesis model for estimating the baseline behavior of a set of neurons, to which we compare our experimental results when quantifying redundancy in groups of neurons.

This model can be refined to take into account the variable information values across the isolated single neurons $I(X_i; S)$. Since in this heterogeneous case the order of the series $(X_1, ..., X_N)$ determines the shape of the curve, we first define the information curve for a predefined ordered set of variables $X_1 \dots X_N$.

$$I(N) = I^{max} \left[1 - \Pi_{i=1}^{n} \left(1 - I(X_i; S)\right)\right] \ . \tag{3.28}$$

For an non-ordered set of variables $\{X_1, ..., X_N\}$, the information curve is defined as the average over all possible orderings. It is interesting to compare this model with the homogeneous model for which I_1 is set to $I_1 = \frac{1}{N} \sum_i I(X_i; S)$. Since for any two numbers $x \times y < \left(\frac{x+y}{2}\right)^2$, the pairs I_1^2 is larger than $< I_1^i I_1^j >$. Similar consideration for higher powers shows that the homogeneous information model yields an underestimation of the curve $I(N)$, i.e. an overestimation of the redundancies.

3.4 Summary

This section has discussed methodological issues in quantifying high-order correlations among small groups of neurons, specifically redundancy and synergy. We defined information theoretic measures, and discussed ways for reliable estimation of these quantities from actual electrophysiological data. The following section describes the application of these methods to neurons in the auditory system.

Chapter 4

Redundancy Reduction in the Auditory Pathway

The previous chapter described quantitative measures of interactions among small groups of neurons. In this chapter these measures are applied to study the neural code in the ascending auditory pathway. Sections 4.1 describes coding of stimulus identity and 4.2 discusses coding of acoustic features .

The current chapter describes the main empirical result of this dissertation, namely, evidence for a process of redundancy reduction along the core ascending auditory pathway.

4.1 Coding stimulus identity

We applied the methods described in the previous chapter to electrophysiological recordings in the core auditory pathway, quantifying mean redundancies in population of neurons in three processing stations: the inferior colliculus (IC), the medial geniculate body of the thalamus (MGB) and the primary auditory cortex (AI). The current section summarizes our findings.

To first demonstrate the type of responses encountered in our data, Fig. 4.1 plots the mean firing rate of two IC neurons in response to 15 stimuli, compared with the mean firing rate of 2 AI neurons to the same stimuli. While the response profiles of the IC neurons to these stimuli are very similar, the AI neurons responded rather differently. For example, one AI neuron (in blue) responded to all stimuli at a roughly constant level,

except for stimuli number 7, 8 and 9. The other neuron, on the other hand, responded to these stimuli at approximately the same level as to the other stimuli, but responded much more weakly to stimulus number 13. These results are not simply due to differences in the frequency response characteristics of the two pairs, because both had similar BF's (5.5 kHz and 6.1 kHz for the IC neurons, 5.1 for both AI neurons).

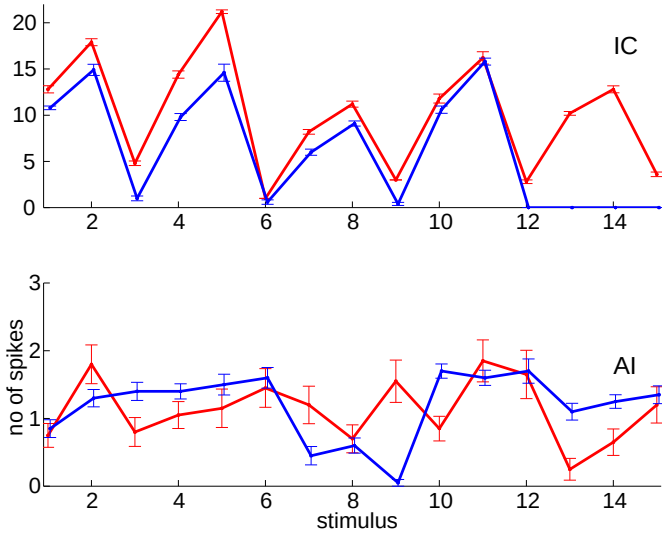

Figure 4.1: Spike counts across the stimulus ensemble for a pair of IC cells (top), (BF's 5.5 and 6.1 kHz) and a pair of AI cells (bottom), (BF's 5.1 both). Error bars denote standard error of the mean spike count, as obtained from 20 repetitions of stimulus presentations.

We quantify pairwise redundancy in neurons from the auditory processing stations. For this purpose we measured the normalized pairs redundancy under conditional independence for all pairs of neurons (as explained in section 3.3.2)

$$Normalized\ Redundancy = \frac{-I(X_1;X_2)}{I(X_1;X_2;S)} = \frac{-I(X_1;X_2)}{I(X_1;S)+I(X_2;S)} \quad (4.1)$$

and plotted its distribution. All information measures were corrected for bias as discussed in chapter 2.

Figure 4.2 plots the distribution of normalized pairwise redundancy with information obtained from spike counts. It reveals a considerable difference

in redundancy level between the IC population on one hand and AI and MGB on the other.

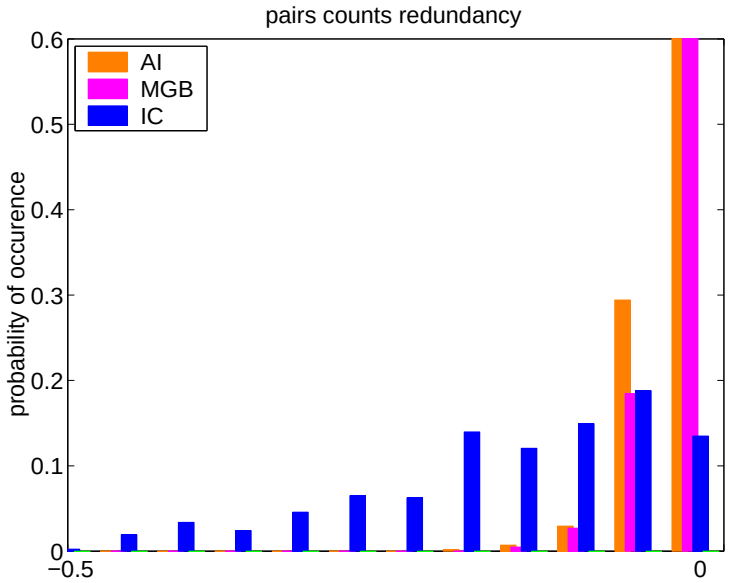

Figure 4.2: Distribution of normalized pairs redundancy $\frac{-I(X_1;X_2)}{I(X_1;S)+I(X_2;S)}$ estimated using spike counts across the population of three brain regions.

Distribution of spike counts was calculated by counting spikes within a window that was optimized for each brain region separately. The window was chosen such that it maximized the mean single-neuron information over the population. The optimal window values were AI: $20-140ms$, MGB: $20-80ms$ and IC: $0-60ms$.

Neurons in A1 and MGB are significantly less redundant than neurons in IC (Fig 3a). The median normalized redundancy in IC was -0.13 (with a median absolute deviation from the median of 0.07), whereas in MGB it was -0.02 (0.015) and in A1 -0.03 (0.015), a highly significant difference.

We further measured normalized triplets redundancies

$$\frac{-I(X_1;X_2;X_3)}{I(X_1;S)+I(X_2;S)+I(X_3;S)} \tag{4.2}$$

in the same population, which showed an even more pronounced difference (Fig. 4.3).

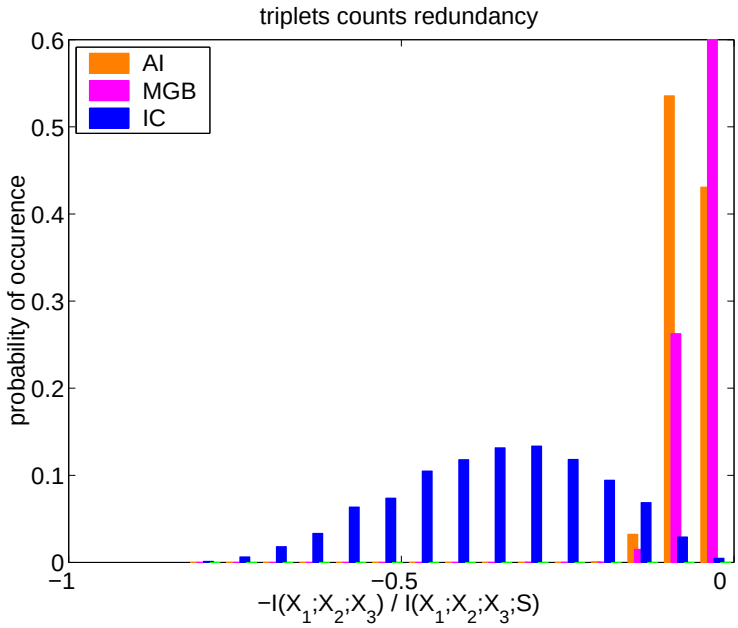

Figure 4.3: Distribution of normalized triplets redundancy across the population three brain regions.

One possible cause for low redundancy in AI compared to IC is the use of a reduced measure, the spike counts. Other statistics of the spike trains could perhaps show comparable redundancies in both areas. We therefore calculated the redundancies using first spike latency as the reduced response measure, and using the direct method (section **??**). Figure 4.4 show that the general picture remained the same, with redundancies in IC substantially larger than in MGB and AI.

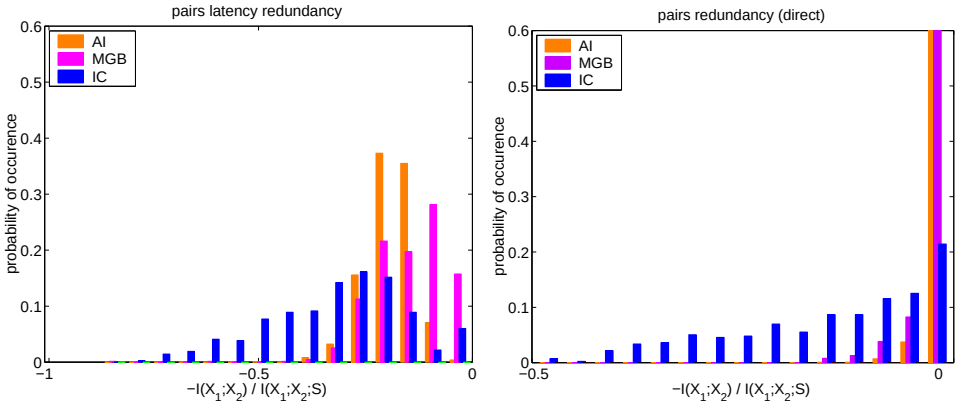

Figure 4.4: Distribution of normalized pairs redundancy across the population of three brain regions, as estimated using first spike latency and the direct method.

To demonstrate that these differences in redundancies are not caused by our normalization procedure, Figure 4.5 shows the distributions of the non normalized redundancy measure under the same conditions, showing that essentially the same results are obtained.

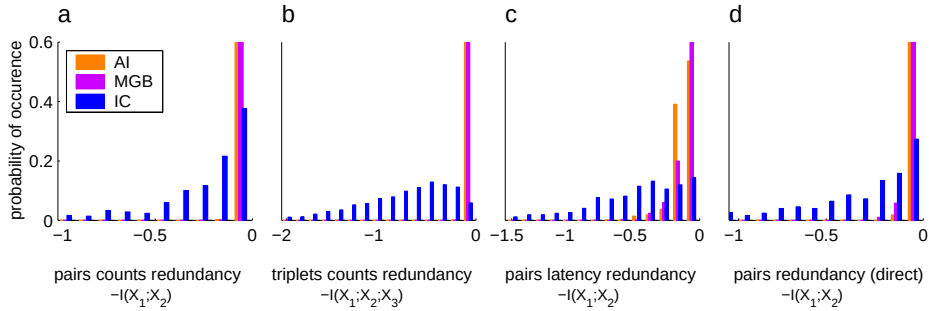

Figure 4.5: Distribution of non-normalized redundancies (A) pairs counts (B) triplets counts. (C) pairs, latencies (D) pairs, direct method .

As summarized in Table 4.6, in all these measures, the IC population of neurons had a considerably larger redundancy than lower level neurons. The following two subsections further refine the characteristics of this redundancy.

method	IC	MGB	AI
counts, pairs	-0.816±0.585	-0.038±0.026	-0.027±0.026
counts, triplets	-0.307±0.094	-0.061±0.025	-0.043±0.025
pairs, latency	-0.332±0.178	-0.186±0.055	-0.142±0.086
pairs, direct	-0.334±0.179	-0.005±0.013	-0.001±0.026

Figure 4.6: Mean and standard deviation of redundancy in four different coding schemes and three brain regions.

4.1.1 Validating the conditional independence approximation

The results of the previous section were obtained using the stimulus conditioned approximation for estimating the joint distributions of pairs and triplets of cells. As discussed in section 3.3.1, this approximation may effect the estimated level of redundancy. We now turn to estimating the validity of this approximation. We measured redundancy and synergy for a smaller number of cells that were recorded simultaneously in AI (total of 9 pairs from 15 cells) and MGB (total of 43 pairs from 29 cells).

Figure 4.7 plots the bias-corrected normalized redundancy obtained from these cells both under the conditional independence approximation and without it. It shows that using the actual coupling between the neurons, instead of the conditional independence approximation, increases the estimated redundancy by a factor of 1.74 on average in MGB cells. The corresponding increase in AI cells was 1.23 on average. It should be noted that redundancies observed in IC were about 5 to 15 times larger than those observed in MGB and AI. We conclude that the conditional independence approximation cannot be the sole cause for this difference in redundancies.

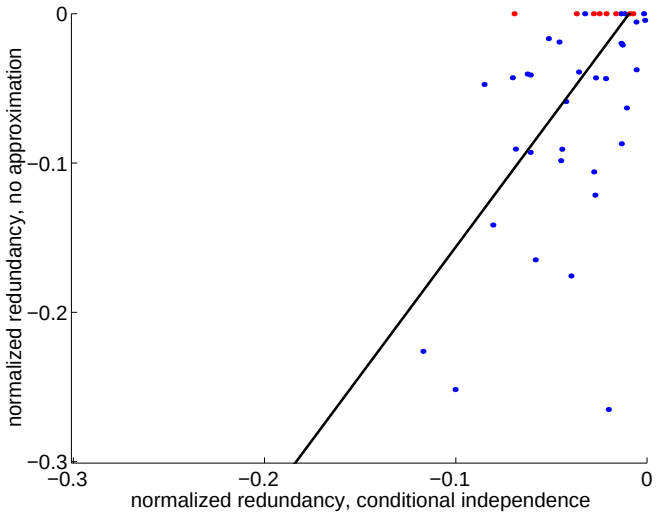

Figure 4.7: Normalized redundancy with and without stimulus conditioned independence approximation in simultaneously recorded MGB neurons. Red points are points that yielded positive redundancy estimate and were clamped to zero.

As explained in section 3.2.1, a frequently used measure of synergy and redundancy in the literature is the difference between the information conveyed by two cells together and that conveyed by the two cells considered individually, $SR_{pairs} = I(X_1, X_2; S) - I(X_1; S) - I(X_2; S)$, which can be rewritten as $I(X_1; X_2|S) - I(X_1; X_2)$.

Figure 4.8 compares this synergy-redundancy index with our redundancy index $-I(X_1; X_2)$ for MGB cells, both normalized by $[I(X; S) + I(Y; S)]$. It shows that in MGB, these two indices are correlated with a linear regression slope of 1.99, and that a linear relation of $SR = -1.0 * I(X_1; X_2)$is within the 95 percent confidence interval of the slope (b $\in$ [0.85, 3.4]). The reason for this correlation is that the synergy in our data is relatively weak (mean synergy -0.02, with a standard deviation of 0.28), and is counterbalanced by an increase in the redundancy when the conditional independence approximation is not used (Fig. 4.7). Similar results were obtained for AI cells (linear regression curve SR=-0.8 I+0.2, mean synergy 0.039 with a standard deviation of 0.12).

We are particularly interested in the redundancies induced by similar frequency sensitivity, and therefore used $-I(X_1; X_2)$ under stimulus-

conditioned independence as our measure of redundancy. The result shown in Fig. 4.8 suggests that the same conclusions would be reached by using SR with the actually measured coupling between the responses. Our measure has two major advantages. First, it can be evaluated reliably for larger sets of neurons (e.g. in Fig. 4, up to 19 neurons). Secondly, it can be evaluated for neurons that have not been simultaneously measured, thus allowing for the use of a larger fraction of the recorded neurons.

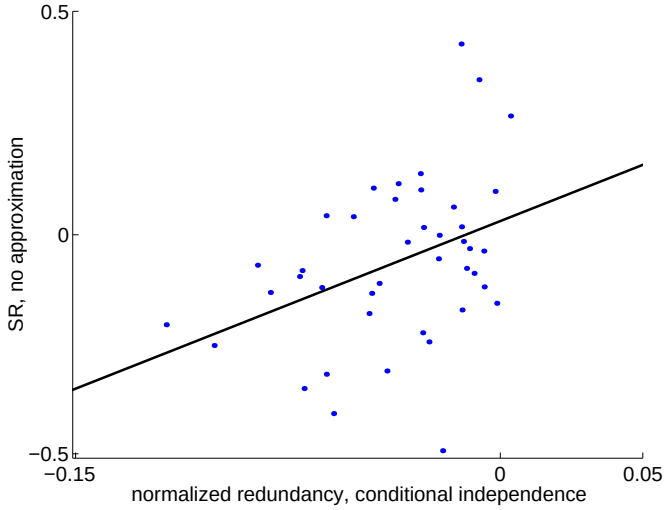

Figure 4.8: Redundancy under conditional independence approximation vs. the standard synergy-redundancy measure in simultaneously recorded MGB neurons.

4.1.2 Redundancy and spectral sensitivity

Neurons in the auditory system are traditionally characterized by their spectral sensitivity. Specifically they are often characterized by the frequency to which they are most responsive, their "best frequency" (BF). A natural question when characterizing redundancy is thus the relation between redundancy levels and spectral sensitivity.

To address this question, we studied the relation between the redundancy of each pair of neurons and their BF's. Figure 4.9 plots the normalized redundancy for each pair of AI neurons, while neurons were ordered by their BF. As a baseline for comparison we used the set of auditory nerve fibers

(ANF) model neurons described in Chapter 1. The set of ANF neurons had the same set of BF's as the AI neurons. Their responses to the same set of stimuli were computed and the redundancy was estimated exactly the same way as for the AI neurons.

Figure 4.9B depicts the normalized redundancy for all pairs of ANF neurons. High redundancy (negative values) is observed along the diagonal, in particular in those frequency bands that have high energy in the stimulus set. Figure 4.9A depicts normalized redundancy values for the AI neurons

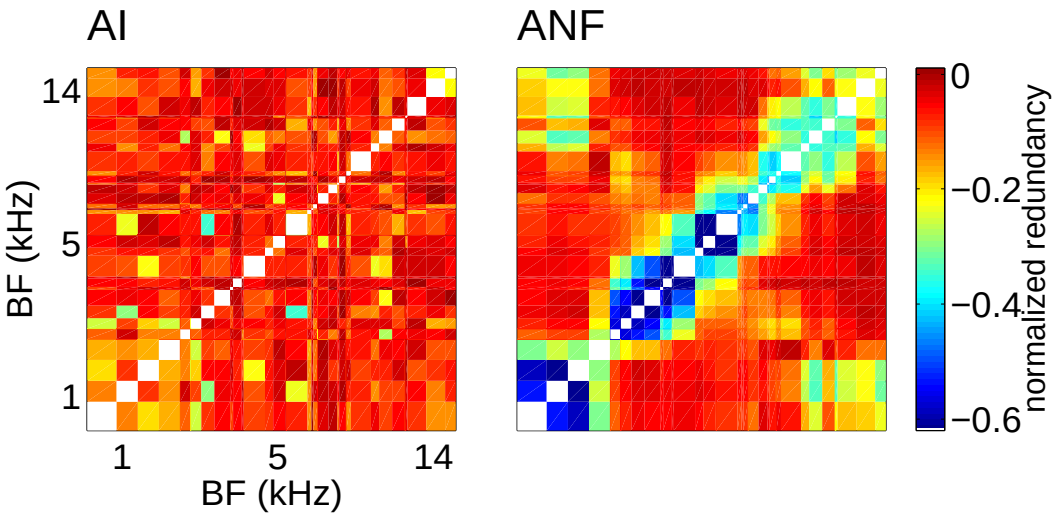

Figure 4.9: Normalized redundancy of pairs of cells ordered by their BF's.

Figure 4.10 plots the normalized redundancy for each pair as a function of the difference in BF's (in log scale), for each of the four brain regions. A strong correlation between BF difference and redundancy level is apparent for both IC and ANF-model neurons, but not for MGB and AI neurons.

In order to quantify this effect, we measured the correlation between the BF difference (in log scale) and the normalized redundancy level. In order to take into account the fact that the points in the scatter plot are the results of pair-wise comparisons, and thus are not independent, we correlated redundancy values against BF difference separately for each unit against all other units. We then tested the hypothesis that the set of correlation coef-

ficients is significantly different from zero. After correcting for the number of samples,

both AI and MGB populations had non significant correlations of redundancy with BF (p¿0.01). On the other hand, in IC and in the ANF simulations, pairs of neurons showed significant correlations of redundancy with BF (IC: $p < 10^{-6}$;ANF: $p < 10^{-8}$.)

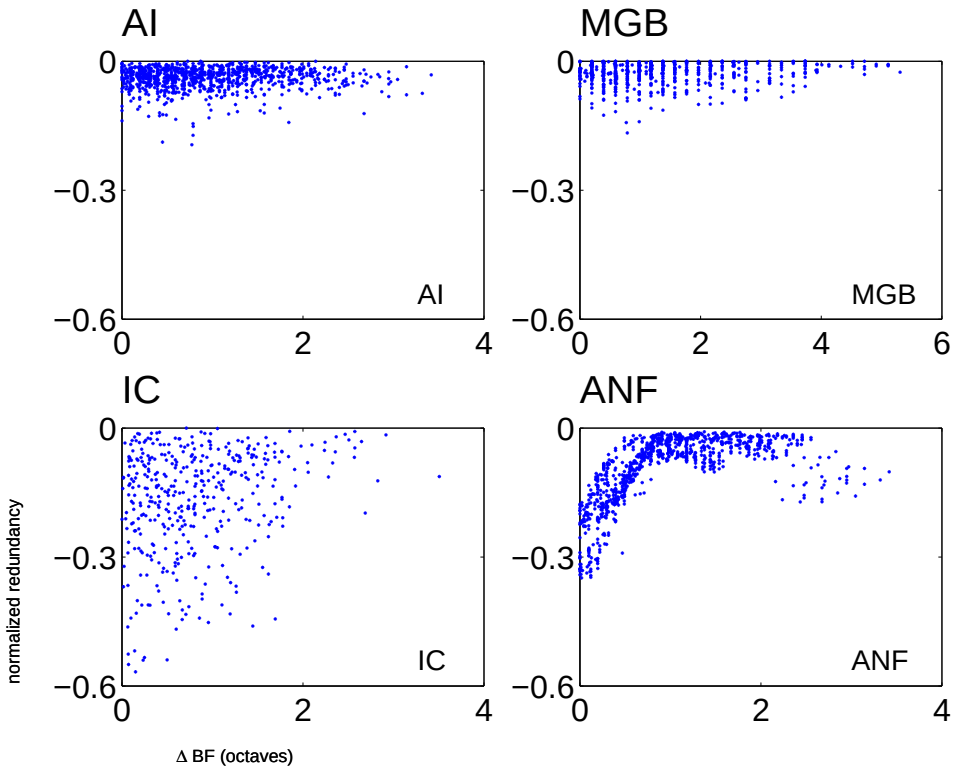

Figure 4.10: Normalized redundancy of pairs of cells as a function of their BF's difference.

The importance of these results lies in the fact that they provide strong hints about the differences in the nature of the neural code in the different stations. IC neurons are more redundant, and this redundancy is related to spectral sensitivity. More specifically, pairs of neurons that show high redundancy, and so convey similar information about the identity of the sounds tend to have similar BF. The reverse is not necessarily true, in that

there are pairs of IC neurons with the same BF that have low redundancy. These are expected, since neurons in IC with the same BF may differ in other aspects of their sensitivity to the physical structure of sounds; e.g. by having different temporal sensitivities [Casseday *et al.*, 2002].

The behavior of AI neurons is inherently different. First, AI neurons are far less redundant. In addition, even the most redundant AI pairs have larger BF difference than the most redundant pairs in IC. The complete lack of redundancy in AI neurons, even among those with similar BF, suggests that different neurons in AI tend to convey information about different aspects of the stimuli, since their redundancy cannot be accounted for by mere spectral sensitivity.

4.1.3 Redundancy and physical cell locations

The previous section characterized the relation between redundancy and cells spectral sensitivity, specifically the distance between cells' BF's. Another possible organization principle underlying redundancy may be revealed in the anatomical organization of the recorded cells. Unfortunately, the 3-dimensional locations of the recorded neurons were not available in our data. However, anatomical distances can be very coarsely estimated by separating the recorded cells into three groups according to their recording type:

- **Same penetration**: Cells that were recorded within a single penetration but on different electrodes, which must be less than 1 mm apart in the cortex. In the MGB, such neurons are less than 1 mm apart in the medio-lateral and antero-posterior axes, but could be more distant in the dorso-medial axis.

- **Same animal**: Cells that were recorded during the same experiment but from different penetration are usually up to a few mm from each other.

- **Different animals**: The distance between cells recorded from different animals is not defined. Since neural organization in different animals could be considerably different, we treat cell pairs in this group as if located at a large distance from each other.

Figure 4.11 plots the normalized redundancy in each of these four groups for MGB cells [1]. Error bars denote the standard errors of the mean of each group. For comparison the mean redundancy in IC is also plotted, showing that MGB redundancy is significantly smaller than IC in all groups. A one-way ANOVA suggests that the difference between the groups is significant ($p < 0.01$) Even more interestingly, the three groups have a monotonically decreasing order of redundancy levels as expected. This means that neurons that are physically near to each other also tend to be more redundant.

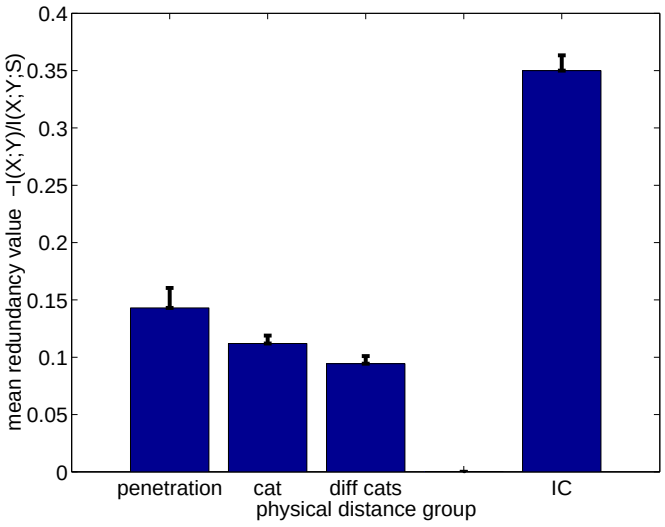

Figure 4.11: Mean normalized redundancy in MGB cell-pairs grouped according to physical distance type, and in IC. Error bars denote SEM of each group.

These results are in agreement with the idea that there is some functional segregation of neurons in the thalamus. Such a functional segregation is also expected in the cortex [Powell and Mountcastle, 1959, Hubel and Wiesel, 1962, Calvin, 1995]. There is ample evidence for the presence of functional clustering and gradients in the Thalamus, both in the visual and auditory modalities. When combined with the results of the previous subsection, it suggests that MGB neurons are physically organized according to their functional sensitivity, but in addition to sensitivity to spectral content of

[1] Unfortunately, for AI and IC recordings there are very few pairs in the first three groups, rendering this analysis unproductive for these data.

sounds, additional acoustic components are used for thalamic organization. Furthermore, the local clustering of functional properties in the MGB is much weaker than in the IC, as suggested by the much lower redundancy values in the MGB, even among neurons recorded in the same penetration.

4.2 Coding acoustics

Neurons in the auditory system are usually analyzed in terms of their spectro-temporal selectivity (e.g. [DeCharms *et al.*, 1998, Schnupp *et al.*, 2001]), i.e. the patterns of stimulus energy across frequency and time to which the neuron is selective. In order to provide insight into the nature of the reduced redundancy among MGB and AI neurons, we investigated the information that single spikes convey about short-term spectro-temporal structures in our stimuli.

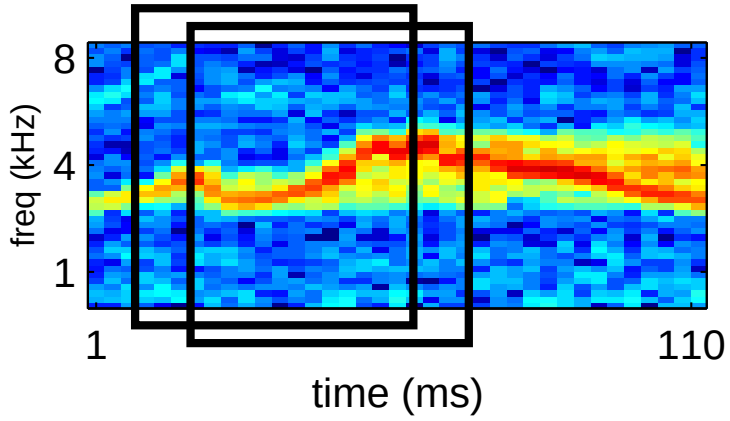

Figure 4.12: Illustration of the way overlapping segments of the bird chirp are created.

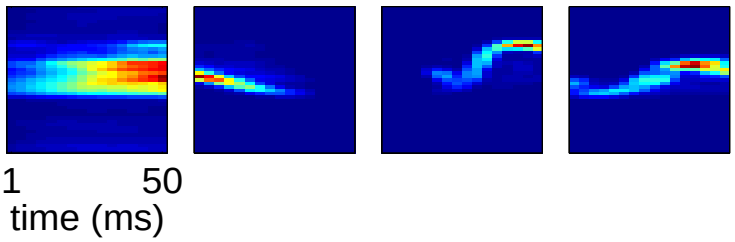

Figure 4.13: Four examples of the means of segment clusters after segments were grouped into 32 clusters.

To this end, the acoustic signal was cut into 50 ms segments with 49 ms overlaps, and the spectrogram of each segment was computed using a

Hanning window with a length of 5.8 ms and a 2.8 ms overlap (Fig. 4.12). This set consisted of about 1500 segments. In order to reduce the dimensionality of the data, we aim to cluster the segments based on their physical properties, namely the spectro temporal energy pattern. The appropriate distance measure is however not known a-priori. We parametrized a family of metrics between segments by passing each segment through a sub-linear transformation $x \rightarrow x^{\alpha}$ in a pixel-wise manner. For α values smaller than 1, this transformation has the effect of emphasizing weak "pixels" as compared to pixels with high energy. The value of α was chosen to maximize the information conveyed by the sample of AI neurons, yielding the value of 0.5. The segments were then clustered using K-means into 32 representatives using a dot product metric operated on the transformed spectrograms. The mean of each of these clusters was calculated (Fig. 4.13).

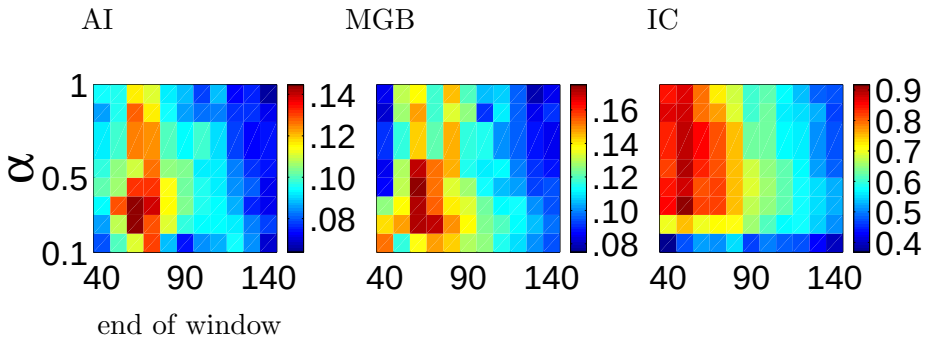

Figure 4.14: Total mutual information obtained from single spikes about acoustics, as a function of the power α and the end point of the temporal window (after stimulus onset). Start time of the window in this figure was taken to be 20 ms after stimulus onset for all three brain regions.

In order to estimate the information that neurons convey about the acoustic segments, the joint probability of spikes and representative segments was estimated. This was achieved by counting combinations of multi-cell spikes patterns, where the responses of a single neuron were considered to be 1 or 0, according to the presence or absence of a spike just following the appearance of a segment from a given cluster. Because of the simplicity of this response set, it was possible to create a joint distribution between clusters and combinations of spike patterns consisting of the responses of multi-

ple neurons, again coupled under the conditional independence assumption.

We then calculated the MI between the spikes evoked in groups of cells and the stimulus clusters immediately preceding the spikes. IC spikes from single cells provided on average 9.2 times more the information than MGB cells, and 7.6 times more than AI cells about the identity of the preceding acoustic segment (Fig. 4.2). This suggests that in contrast with IC cells, AI and MGB neurons poorly discriminate between stimuli grouped on the basis of spectro-temporal energy structure. The underlying reason is the high sensitivity of AI neurons to small perturbations in the stimuli, as illustrated in Fig. 1B. This is in contradistinction to the way they code the identity of the stimuli, abbot which they convey almost half the information, compared to the same IC cells (Fig. 2.8 on chapter 2). Therefore, AI neurons distinguish between complex stimuli much better than would be expected based on the information they confer about the acoustic structure.

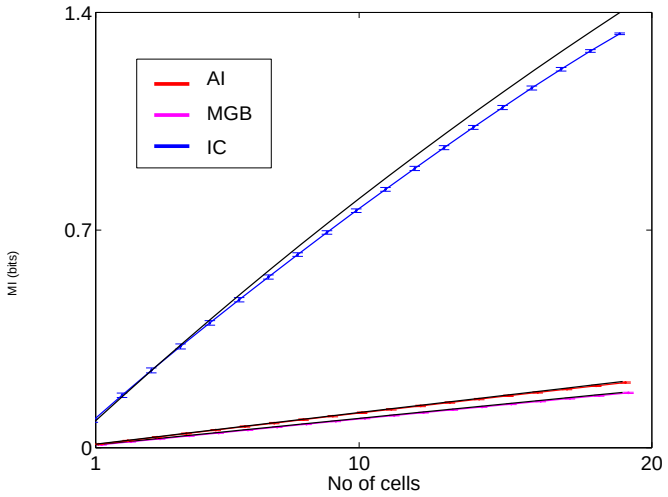

Figure 4.15: Information conveyed by single spikes about acoustic segments as a function of the number of neurons: IC-blue, MGB-magenta (partially covered by AI), A1-red. The black curves denote the expected information obtained from independent neurons. Error bars designate the standard error of the mean MI for several subsets of the same size. For each set size, analysis was repeated for 20 subsets and for 5 different randomization seeds.

To estimate redundancy, MI from a group of cells was again compared with the MI expected from independent cells, as explained in section 3.2.6.

The normalized redundancy index (Fig. 4.2) reveals a significantly larger deviation from independence for IC cells than for MGB and cortical cells. These results suggest that in the coding of short-term spectro-temporal structure, as in the coding of stimulus identity, neural representations change along the ascending auditory system in a way that reduces redundancy among neurons.

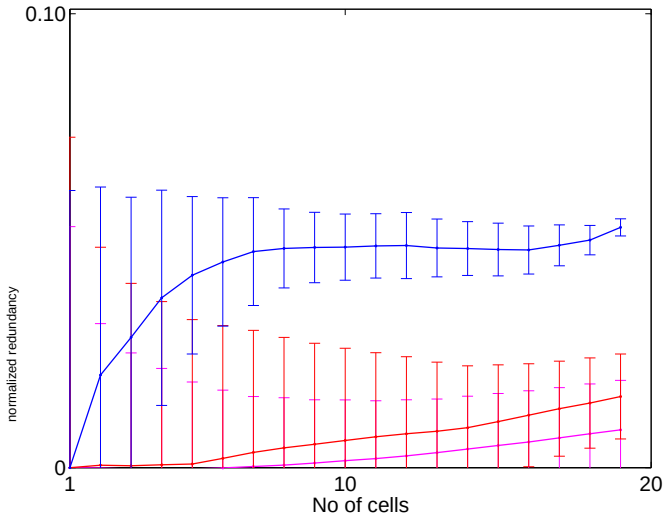

Figure 4.16: Normalized redundancy: the difference between the obtained information and the baseline of expected information from independent cells, normalized by the baseline.

4.3 Summary

This chapter has investigated the interactions among small groups of auditory neurons, and their relation to coding of acoustic stimuli. The unique setup of our dataset, namely, responses to the same set of stimuli in a number of auditory processing stations, allowed us to compare these interactions in different brain regions. We developed rigorous quantitative measures of information redundancy among sets of neuronal responses and applied them to the electrophysiological data.

Our main results are threefold. First, we showed that small groups of IC cells tend to be more redundant in the information they convey about

the stimulus identity than AI and MGB cells. In other words, cells of higher regions in the processing hierarchy code stimuli in a more informationally independent manner. Secondly, we showed that this redundancy is significantly correlated with the BFs the IC cells but not with the BFs of AI or MGB cells. This means that frequency characterization poorly captures the type of processing neurons in MGB and AI perform. Finally, AI and MGB cells convey an order of magnitude less information about the spectro-temporal structure of the stimuli as compared with IC neurons. This suggest that neurons in MGB and AI succeed in coding the identity of the stimuli but without coding well the precise acoustical structures in it.

The low redundancy in AI and MGB, and the lack of correlation of this redundancy with the BFs of the cells, has strong implications, which go far beyond the assertion that BF responses are predictive for complex sounds. The reason is that current accepted methods of characterizing AI neurons, primarily the spectro temporal receptive field (STRF), imply redundancy between neurons that share spectro temporal characteristics. Although we find such redundant neurons in IC we do not find them in MGB or AI, demonstrating that STRF characterization in AI misses crucial aspects of neuronal coding even for stimuli as simple as those used in our work.

Bibliography

[Abeles *et al.*, 1993] M. Abeles, H. Bergmann, H. Margalit, and E. Vaadia. Spatiotemporal firing patterns in the frontal cortex of behaving monkies. *The Journal of Neurophysiology*, 70:1629–1638, 1993.

[Bialek *et al.*, 1991] W. Bialek, F. Rieke, R.R. deRuyter van Steveninck, and D. Warland. Reading a neural code. *Science*, 252:1854–1857, 1991.

[Brenner *et al.*, 2000] N. Brenner, S.P. Strong, R. Koberle, R. de Ruyter van Steveninck, and W. Bialek. Synergy in a neural code. *Neural Computation*, 13(7):1531–1552, 2000.

[Calvin, 1995] W.H. Calvin. Cortical columns, modules, and Hebbian cell assemblies. In Michael A. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, pages 269–272. MIT Press, Cambridge, MA, 1995.

[Casseday *et al.*, 2002] JH. Casseday, T. Fremouw, and E. Covey. The inferior colliculus. In Donata Oertel, Richard R. Fay, and A. N. Popper, editors, *Integrative Functions in the Mammalian Auditory Pathway*, volume 15 of *Handbook in Auditory Research*, pages 238–318. Springer-Verlag, New York, 2002.

[Cover and Thomas, 1991] T.M. Cover and J.A. Thomas. *The elements of information theory*. Plenum Press, New York, 1991.

[Csiszar and J.Korner, 1997] I. Csiszar and J.Korner. *Information theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, New York, 2nd edition, 1997.

[Dan *et al.*, 1998] Y. Dan, J.M. Alonso, W.M. Usrey, and R.C. Reid. Coding of visual information by precisely correlated spikes in the LGN. *Nature Neuroscience*, 1:501–507, 1998.

[DeCharms *et al.*, 1998] R.C. DeCharms, D.T. Blake, and M.M. Merzenick. Optimizing sound features for cortical neurons. *Science*, 280(5368):1439–1444, 1998.

[Gat and Tishby, 1999] I. Gat and N. Tishby. Synergy and redundancy among brain cells of behaving monkeys. In M.S. Kearns, S.A. Solla, and D.A. Cohn, editors, *Advances in Neural Information Processing Systems*, volume 11, Cambridge, MA, 1999. MIT Press.

[Gawne and Richmond, 1993] T.J. Gawne and B.J. Richmond. How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience*, 13(7):2758–2771, 1993.

[Gochin *et al.*, 1994] P.M. Gochin, M. Colombo, G. A. Dorfman, G.L. Gerstein, and C.G. Gross. Neural ensemble coding in inferior temporal cortex. *The Journal of Neurophysiology*, 71:2325–2337, 1994.

[Hopfield, 1995] J.J. Hopfield. Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376(6535):33–36, 1995.

[Hubel and Wiesel, 1962] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160:106–154, 1962.

[Jordan, 1998] M. Jordan, editor. *Learning in graphical models*. MIT press, Cambridge, MA, 1998.

[Meister *et al.*, 1995] M. Meister, L. Lagnado, and D. Baylor. Concerted signaling by retinal ganglion cells. *Science*, 270:1207–1210, 1995.

[Meister, 1996] M. Meister. Multi neuronal coding in retinal signaling. *Proceedings of the National Academy of Sciences of the United States of America*, 93:609–614, 1996.

[Nirenberg *et al.*, 2001] S. Nirenberg, S.M. Carcieri, A.L. Jacobs, and P.E. Latham. Retinal ganglion cells act largely as independent encoders. *Nature*, 411:698–701, 2001.

[Panzeri *et al.*, 1999] S. Panzeri, S. R. Schultz, A. Treves, and E.T. Rolls. Correlations and the encoding of information in the nervous system. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 266:1001–1012, 1999.

[Pearl, 1988] J. Pearl. *Probabilistic inference in intelligent systems*. Morgan Kaufman, San Mateo, CA, 1988.

[Powell and Mountcastle, 1959] T.P.S. Powell and V.B. Mountcastle. Some aspects of the functional organization of the cortex of the postcentral gyrus of the monkey: A correlation of findings obtained in a single unit analysis with cytoarchitecture. *Bulletin of the Johns Hopkins Hospital*, 105:133–162, 1959.

[Reich *et al.*, 2001] D.S. Reich, F. Mechler, and J.D. Victor. Independent and redundant information in nearby cortical neurons. *Science*, 294:2566–2568, 2001.

[Rolls *et al.*, 1997] E. T. Rolls, A. Treves, and M.J. Tovee. The representational capacity of the distributed encoding of information provided by populations of neurons in primate termpoal visual cortex. *Experimental Brain Research*, 114:149–162, 1997.

[Samengo and Treves, 2000] I. Samengo and A. Treves. Representational capacity of a set of independent neurons. *Physical Reviews E*, 63:1–14, 2000.

[Samengo, 2001] I. Samengo. Independent neurons representing a finite set of stimuli: Dependence of the mutual information on the number of units sampled. *Network: Computation in Neural Systems*, 12:21–31, 2001.

[Schnupp *et al.*, 2001] J.W.H. Schnupp, T.D. Mrsic-Flogel, and A.J. King. Linear processing of spatial cues in primary auditory cortex. *Nature*, 414:200–204, 2001.

[Shafer and Pearl, 1990] G.R. Shafer and J. Pearl. *Reading in uncertain reasoning*. Morgan Kaufman, San Mateo, CA, 1990.

[Singer and Gray, 1995] W. Singer and C.M. Gray. Visual feature integration and the temporal correlation hypothesis. *Annual Reviews in Neuroscience*, 18:555–586, 1995.

[Vaadia *et al.*, 1995] E. Vaadia, I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A. Aertsen. Dynamics of neural interactions in monkey cortex in relation to behavioral events. *Nature*, 373:515–518, 1995.

[Warland *et al.*, 1997] DK. Warland, P. Reinagel, and M. Meister. Decoding visual information from a population of retinal ganglion cells. *The Journal of Neurophysiology*, 78:2336–2350, 1997.

[Yeung, 1997] R.W. Yeung. A framework for linear information inequalities. *IEEE Transactions on Information Theory*, 43:1924–1934, 1997.