Chapter 1

Introduction

1.1 Computation in the brain

1.1.1 Introduction

What do we mean when we say that the brain computes?

It is not easy to explain to the educated layman what it means that the brain computes. Often, the immediate source of confusion is that the term does not refer to a person performing calculations in his head, but rather to the operations of small circuits of neurons in his brain. The clearest way to think about it is to view computations as mappings, or (possibly high dimensional) functions. By this view, the mapping of addition is simply to map the elements *two* and *two* to an element *four*. This mapping also maps *three* and *one* to the same element *four*. In this context the theory of computation is about studying the ways in which simple mappings can be combined to create complex ones. More complicated functions can map a large set of real numbers into a smaller set that extracts important invariances; for example, by mapping arrays of gray level pixels into a small set of familiar faces, or arrays of sound-pressure levels into a set of comprehensible words. Such mappings can result from the computations performed by our sensory organs, and this dissertation centers on understanding the rules that govern them.

Can we understand how the brain computes?

The extreme difficulty in understanding such complex mappings, is only realized when the relevant quantities are stated. The influx of sensory information to a single human retina is detected by an array of millions of receptors, each capable of telling the difference between hundreds of gray levels, and having time-constants that allow them to detect dozens of new signals in a second. This input is then processed by hundreds of millions of other neurons, many of them interact with each other in complex ways that are constantly changed by the very same inputs we wish to investigate. This architecture is therefore capable of implementing extremely complex maps.

With this gigantic influx, the experimental tools available today are devastatingly weak. The current work uses electrophysiological recordings from small groups of isolated neurons. The data analyzed here were collected from about one hundred neurons only, but required several years of dedicated work done by my collaborators.

With this mismatch between the complexity of the problem and the weakness of the tools, how can we hope to obtain a well established understanding of complex neural systems? The answer lies in the hope that the system adheres to regularities and similarities that simplify the mapping it implements. For example, since neighboring neurons across the neural epithelium are exposed to similar inputs, their functions are expected to share similar properties. This suggests that averages over localized groups of neurons can improve signal to noise issues and allows for extracting coarse maps. Alternatively, developmental and evolutionary considerations can pose additional constraints on the type of maps and computations we may find.

Finally, and this is the approach taken in this thesis, there is hope that these mapping obey some generic *design principles* that guide the type of computation the neural system performs. If such principles exist, we should be able to characterize them more easily than the complex maps themselves, since they will be reflected in multiple subsystems, areas and forms. Moreover, they are expected to embody the functional properties of the neural circuits, which is our ultimate goal in understanding the neural system.

1.1.2 Design principles for sensory computation

The search for design principles that govern the processing performed by sensory systems, was boosted by the appearance of Shannon's information theory in the early 50's. Analogies between sensory systems and communication channels were suggested [Attneave, 1954, Miller, 1956], laying the ground for postulating optimization principles for neural circuits. Although several researchers discussed generic design principles that could underlie sensory processing (e.g. [Barlow, 1959b, Uttley, 1970, Linsker, 1988, Atick and Redlich, 1990, Atick, 1992, Ukrainec and Haykin, 1996, Becker, 1996], and see also chapter 10 in [Haykin, 1999]), I focus here on *Information Maximization* and *Redundancy Reduction*.

Information maximization

The information maximization principle (InfoMax) put forward by Linsker [Linsker, 1988, Linsker, 1989], suggests that a neural network should tune its circuits to maximize the mutual information between its outputs and inputs. Since the network usually has some prefixed architecture, this amount to a constrained optimization problem for any given set of inputs. This approach was used in [Linsker, 1992] to devise a learning rule for single linear neurons receiving Gaussian inputs. It was extended in [Linsker, 1997] to the case of multiple output neurons utilizing local rules only in the form of lateral inhibition.

While Infomax was originally formulated such that the input-output mutual information maximization is the goal of the system, it was extended to other scenarios. Becker and Hinton (1992,1996) presented *Imax*, one of the important variants of Infomax in which the goal of the system is to maximize the information between the outputs of two neighboring neural networks. They showed how this architecture can be used to extract spatially coherent features in simulations of visual processing. Another variant was presented by [Ukrainec and Haykin, 1996], where the goal of the system was the opposite of that of *Imax*. They showed how minimization of mutual information between outputs of neighboring networks extracts spatially incoherent features, and can be usefully applied to the enhancement of radar images. In

[Uttley, 1970] the *Informon* principle was described, where minimization of the input-output mutual information was used as the optimization goal. Such a system becomes discriminatory of the more frequent patterns in the set of signals.

In a paper which is not included in this dissertation [Chechik, 2003], I showed how Infomax can be extended to maximize information between the output of a network and the identity of an input pattern. This setting allows to extract *relevant information* using a simplified learning signal, instead of reproducing the networks inputs. Interestingly the resulting learning rule can be approximated by a spike time dependent plasticity rule.

Redundancy reduction

Redundancies in sensory stimuli were put forward as important for understanding perception since the very early days of information theory [Attneave, 1954]. Indeed these redundancies reflect structures in the inputs that allow the brain to build "working models" of its environments. Barlow's specific hypothesis [Attneave, 1954, Barlow, 1959b, Barlow, 1959a, Barlow, 1961] was that one of the goals of a neural system is to obtain an efficient representation of the sensory inputs, by compressing its inputs to achieve a parsimonious code. During this compression process, statistical redundancies that are abundant in natural data and therefore also characterize the representation at the receptor level, are filtered out such that the neuronal outputs become statistically independent. This principle was hence named *Redundancy Reduction*. The redundancy reduction hypothesis inspired Atick and Redlich (1990), to postulate the principle of *minimum redundancy* as a formal goal for learning in neural networks. Under some conditions [Nadal and Parga, 1994, Nadal *et al.*, 1998] this minimization of redundancy becomes equivalent to maximization of input-output mutual information.

Achieving compressed representations provides several predictions about the nature of the neural code after compression, namely that the number of neurons required is smaller but their firing rates should be higher. The neurophysiological evidence however does not support these predictions: For example, the number of neurons in the lower levels of the visual system is ten times smaller than in the higher ones, and the firing rates in auditory cortex are significantly lower than in the auditory nerve. This suggests that parsimony may not be the primary goal of the system.

Barlow then suggested [Barlow, 2001] that the actual goal of the system is rather *redundancy exploitation*, a process during which the statistical structures in the inputs are removed in a way that reflects the fact that the system used it to identify meaningful objects and structure in the input. These structures are later represented in higher processing levels, a process that again yields a reduction in coding redundancies of higher level elements.

1.2 Information theory

Information theory plays several different roles in the current thesis: both conceptual and methodological. At the methodological level, we use the basic quantities of information theory - such as entropy, mutual information, and redundancy - to quantify properties of the stochastic neural activity. But more importantly, information theory provides a conceptual framework for thinking about design principles of neurally implemented maps. Finally, we also use information theoretic tools to develop methods of unsupervised learning to make sense of the data.

The fundamental concepts of Information theory are reviewed in Appendix B. The reader is referred to [Cover and Thomas, 1991] for a fuller exposition.

1.3 To hear a neural code

In the studies described in this dissertation the main data source were electrophysiological recordings in the auditory system of cats. To understand the findings presented in the main body of the thesis, I now provide a short review of the architecture of this system, both in terms of its gross anatomy and its physiology.

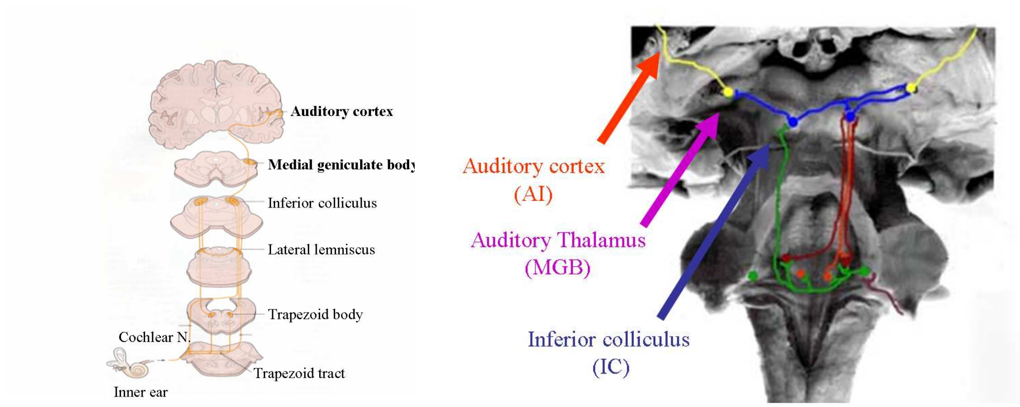

Figure 1.1: **Left.** An illustration of the anatomy of the mammalian auditory system. **Right.** A cross section of a human brain, on which the auditory pathway is marked. The three auditory processing stations analyzed in this work are designated: IC, MGB and AI.

1.3.1 Gross anatomy of the auditory system

This thesis focuses on the core pathway of the auditory system in mammals. This pathway consists of several neural processing stations: the 8th (auditory) nerve, the cochlear nucleus (CN), the superior olivary complex (SOC) and the nuclei of the lateral lemniscus (NLL), the inferior colliculus (IC), the medial geniculate body of the thalamus (MGB), and the primary auditory cortex (AI) [Popper and Fay, 1992]. An illustration of the mammalian auditory system is presented in Figs. 1.1.

In addition to the ascending system, there is also a strong descending information flow, where the major descending pathway projects from the cortex to the thalamus and IC, from IC to lower centers and finally from sub-nuclei of the SOC to the cochlear nucleus and to the cochlea [Spangler and Warr, 1991].

The next subsections briefly review some of the main functional properties of the processing stations of the core pathway, and provide a few examples of raw data later used in the analysis presented in the main chapters of the thesis. Aspects of localization or binaural processing are not discussed here, and the interested reader is referred to [Middlebrooks *et al.*, 2002].

1.3.2 Auditory nerve fibers

Auditory nerve fibers project information from the auditory receptors (the hair cells of the cochlea) into the cochlear nucleus, which is the first processing station of acoustic stimuli[1]. To characterize the spectral sensitivity of an auditory nerve fiber, pure tones at different frequencies and amplitudes are presented to an animal and the response of the fiber is recorded. At every frequency, the minimal sound level that elicits a significant response is recorded, resulting in a frequency tuning curve, an example of which is presented in Fig. 1.2. It shows that the typical frequency tuning curve consists of a fairly narrow frequency band to which the neurons are sensitive. The frequency that has the lowest threshold is called the *best frequency* (BF).

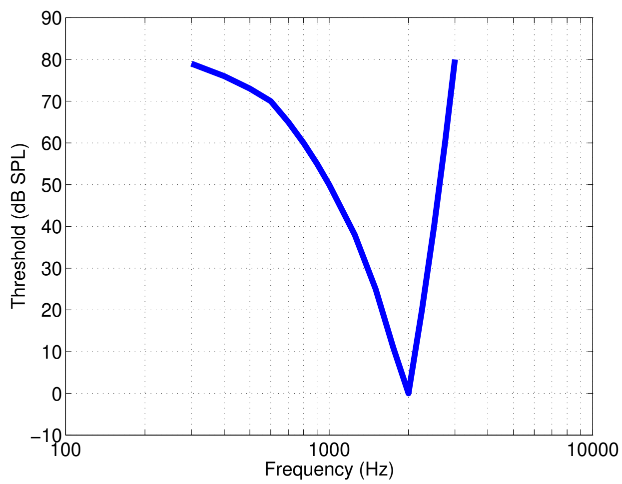

Figure 1.2: **A.** A typical frequency tuning curve of an 8^{th} nerve. It is sensitive to a band of frequencies only few kHz wide. Reproduced from publicly available data

In spite of the presence of strong non-linearities in the responses of auditory nerve fibers, the responses of the population of auditory nerve fibers can be reasonably well described to a first approximation as the output of a band-pass filter bank. A useful model of the responses of these cells is with Gamma-tone filters, where the BF's of the cells are homogeneously spaces along a logarithmically scaled frequency axis. Figure 1.3 depicts an example

[1]The first synapse in the pathway is between the hair cells and the auditory nerve fibers, and the second synapse is in the CN.

of such a set of filters.

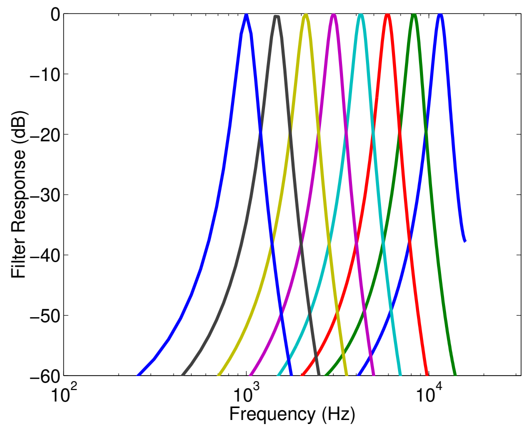

Figure 1.3: The filter coefficients for a bank of Gamma-tone filters. Taken from the auditory tool box by [Slaney, 1998]. Filters were designed by Patterson and Holdworth for simulating the cochlea.

Most interestingly, when auditory nerve fibers are probed with complex sounds such as bird chirps, their response profile can be well explained by their frequency profiles. This is demonstrated in the activity of a neuron from the ventral cochlear nucleus in Fig. 1.4. Whenever the stimulus (middle panel) contains energy in the range of frequencies within the neuron's tuning curve (left panel) as depicted with black horizontal lines, a significant rise in firing rate is observed (lower panel).

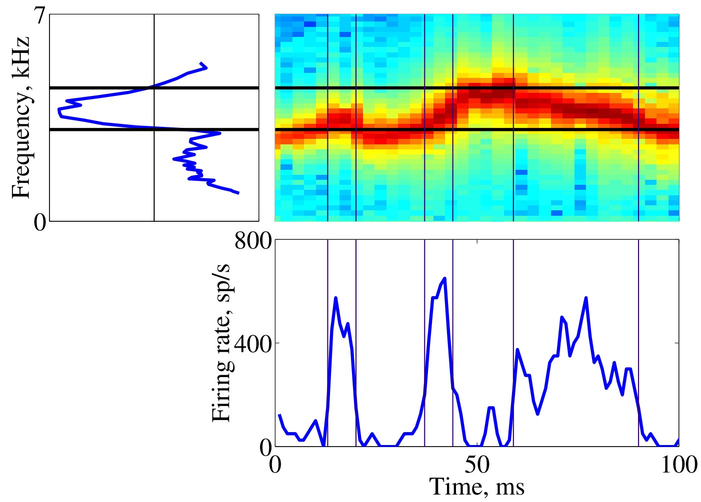

Figure 1.4: Responses of a primary like neuron in the ventral cochlear nucleus, whose behavior is also typical of an auditory nerve fiber. **Left:** A tuning curve. The blue line denotes the minimal level at which a significant response is observed. **Right:** A spectrogram of a bird chirp. Horizontal lines depict the range of frequencies for which the neuron is sensitive. **Bottom:** Firing rate in responses to the presentation of the bird chirp.

Figure 1.5 shows post stimulus time histograms (PSTH) of a model neuron in our data set, as a response to the presentation of a natural bird chirp. The behavior of this model neuron is similar to the recorded ones, in the sense that a coarse but good prediction of the responses to complex sounds can be obtained from the frequency characterization.

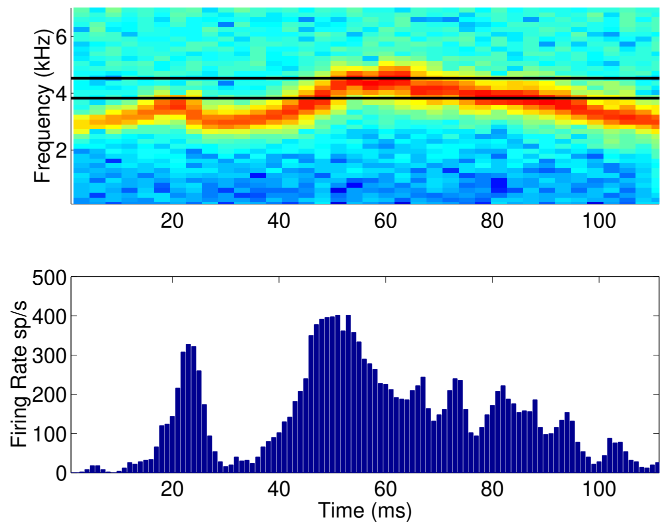

Figure 1.5: Post stimulus time histogram of the responses of a model 8th nerve neuron from our data set, created to have a best frequency at 4.5 kHz.

1.3.3 Inferior colliculus

The inferior colliculus (IC) is an obligatory station of the auditory pathway. All the separate brainstem pathways (from the CN, the SOC and the NLL) terminate in the IC. In addition, the IC receives input from the contralateral IC, descending inputs from the auditory cortex and even somatosensory inputs. As in other auditory areas, IC neurons are frequency sensitive, and exhibit a tonotopic organization within the IC. Moreover the best frequencies of neurons are arranged in an orderly manner, in a way that is fairly well preserved across mammalian species. Most interestingly, the inputs from the multiple origins converge in an arrangement that corresponds to the tonotopic organization of the IC. The IC therefore preserves the same tonotopic map for multiple converging inputs, allowing for complex integration of information within a localized area of the iso-frequency sheet. Not much is known about the organization orthogonal to the frequency gradient, although there is strong evidence for functional gradients related to temporal characteristics, such as best modulation frequency [Schreiner and Langner, 1997] and latency. IC neurons exhibit a rich spectrum of frequency sensitivities, some are sharply tuned to frequencies while some respond to broad-band noise. Some shaping of the frequency tuning is achieved by lateral inhibition [Ehret and Merzenich, 1988] , and some by other mechanisms [Palombi and Caspary, 1996]. Many IC neurons are also selective to temporal structures. Temporal processing in IC include selectivity to sound durations, delays, frequency modulated sounds and more (see section 3.3 in [Casseday *et al.*, 2002] for detailed review). Despite all this, there is still no satisfying description of IC organization in terms of ordered functional maps.

Figure 1.6 presents responses of a typical IC neuron analysed in the current work. When presented with bird chirps, IC response tended to be locked to some features of the stimuli, as indicated by the reliable and precise nature of spikes revealed across repeated stimulus presentations.

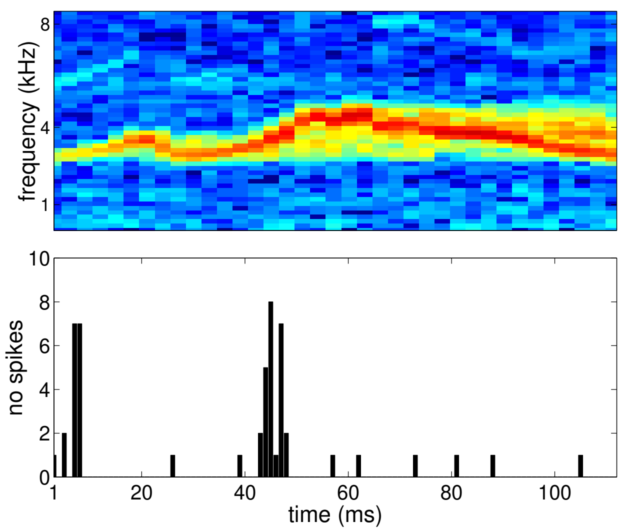

Figure 1.6: Post stimulus time histogram of the responses of an IC neuron from our data set. Notice the tight and precise locking of responses to the stimulus.

1.3.4 Spectro temporal receptive fields in auditory cortex

The auditory cortex is a focus of special interest in this work, since it is the highest processing station we investigated and presumably contains the most complex response properties.

Frequency sensitivity, as characterized with pure tones, reveals that many cortical neurons have a narrow frequency tuning curve, limited dynamic range and often are non monotonic in their responses. Cortical neurons show strong sensitivity to the shape of the tone onset, a dependence that is currently well understood [Heil and Irvine, 1996, Heil, 1997, Fishbach *et al.*, 2001].

When using more complex stimuli, the picture becomes drastically more complex. While FRA characterization could be used to obtain a good description of responses to complex sounds in ANF, this is no longer the case for cortical neurons. The FRA measured by pure tones fails to capture two important aspects of cortical processing: integration across frequencies, and sensitivity to temporal structures [Nelken, 2002].

It was suggested that a better model of cortical responses can be obtained by deriving a spectro temporal receptive field (STRF), an approach that was found useful for characterizing auditory neurons in several systems (e.g. [Aertsen and Johannesma, 1981, Eggermont *et al.*, 1983]). deCharms and colleagues [DeCharms *et al.*, 1998] used short random combinations of pure tones and a spike triggered averaging analysis to obtain the STRF of auditory cortical neurons in monkeys. The resulting receptive fields, demonstrated in Fig. 1.7, show complex dependencies between time and frequency, suggesting that cortical neurons are sensitive to frequency changes, as in FM sweeps. However, this type of analysis is linear in the sense that it averages the energy in spectro-temporal “pixels” while assuming independence between pixels, and it is therefore limited in its ability to capture complex interactions between frequencies and temporal structures.

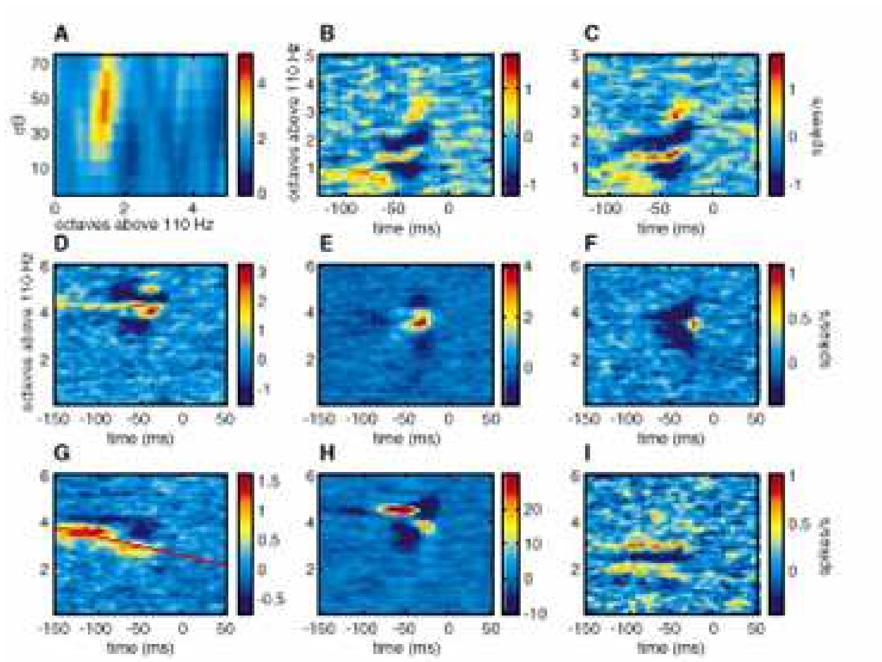

Figure 1.7: **A.** Frequency response area (FRA) of a typical cell. Notice the non monotonic response as a function of level. **B.-I.** spectro temporal receptive fields of different cells. B and C are STRFs estimated from the responses of the neuron in A using different sets of random chords. From [deCharms 1998] .

A striking demonstration of such nonlinear interactions was observed in cortical responses to natural and modified bird chirps [Bar-Yosef *et al.*, 2001]. Bar Yosef and colleagues showed that relatively minor modifications of the stimulus, such as the removal of the background noise from a natural recording, could dramatically alter the responses of cortical neurons. This type of behavior cannot be explained using linear combinations of STRF's. These results are discussed together with the set of stimuli we used, in the next section.

1.3.5 The stimulus set: To hear a mocking bird

In order to study changes in stimulus representation along the processing pathway, one should use a set of stimuli whose processing is not limited to low level processing stations, otherwise, the properties of high level representation will only reflect low level rather than high level processing. Auditory neurons are often characterized by their spectro temporal properties, however, since the exact features for which cortical neurons are sensitive to are still unkown, we chose to use here a stimulus set, that contains several natural stimuli, which contain rich structures in terms of frequency spectrum and modulation. In addition, we added several variants of these stimuli that share some of the spectro temporal structures that appear in the natural stimuli. This set of stimuli is expected to yield redundant representations at the auditory periphery, and is therefore suitable for the investigation of informational redundancy

The stimulus set used here was created by O. Bar-Yosef and I. Nelken and is described in details in [Bar-Yosef *et al.*, 2001]. It is based on natural recordings of isolated bird chirps, whose sound wave and spectrograms are depicted in Fig. 1.8.

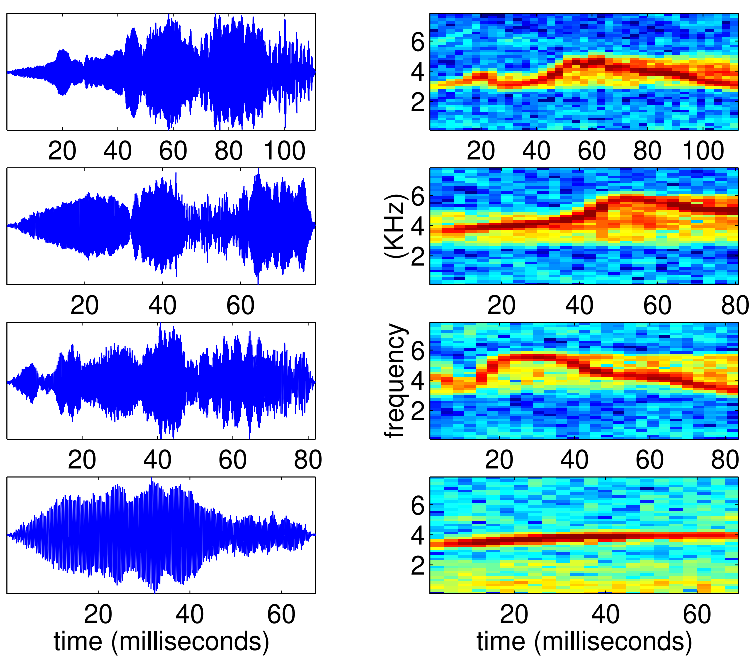

Figure 1.8: Four natural recordings of bird chirps. For each chirp, the left panel shows its sound wave and the right panel its spectrogram.

Each natural recording was then separated into two components: the *main* chirp and the background. The background was further separated into the *echo* component, and the rest of the signal, termed *noise*. These components were then combined into several configurations (*main+echo*, *main + background*). In addition, an artificial stimulus that follows the main FM sweep of the chirp was also created (termed *artificial*). The variants based on the first bird chirp are depicted in Fig. 1.9

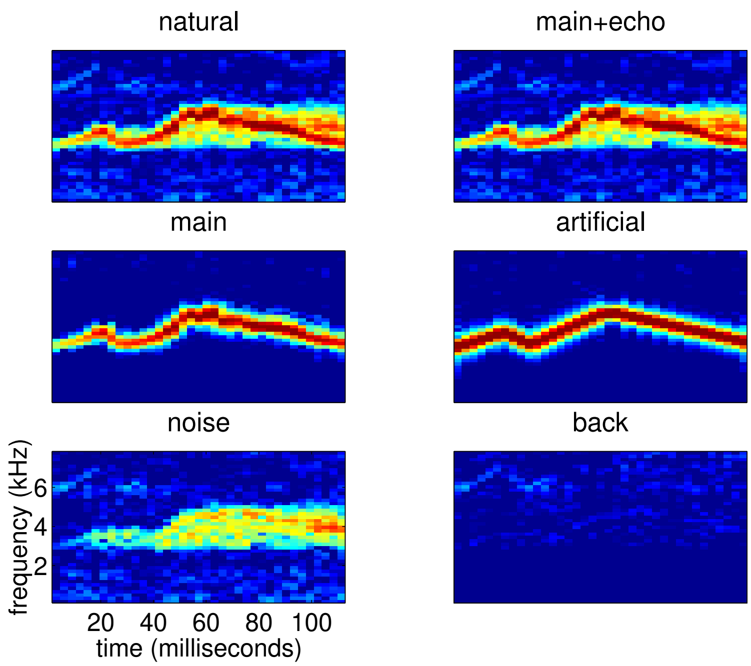

Figure 1.9: Six different variants created from a single natural bird chirp (upper panel in the previous figure) .

In some of the analyses, 32 different stimuli, based on 8 variants of 4 different bird chirps were used. In others, 15 stimuli, based on 5 variants of 3 bird chirps were used. These 15 stimuli are plot below

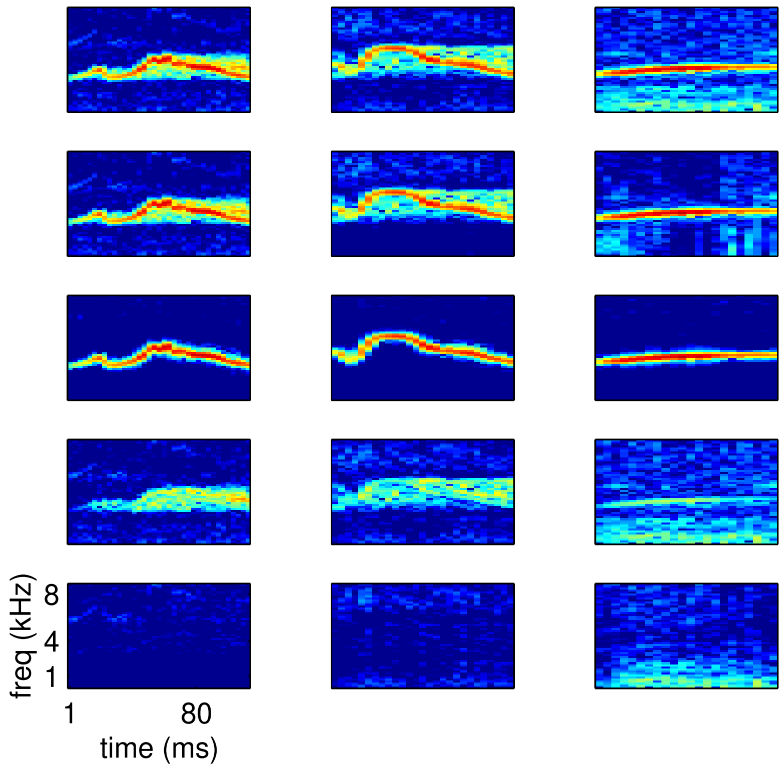

Figure 1.10: A set of 15 stimuli created from three different bird chirps..

1.3.6 The experimental setup

The electrphysiological recordings that provided the data that are analyzed in this work were performed in two laboratories. First is the laboratory of Prof. Eric Young at Johns Hopkins University, Boltimore, where electrophysiological recordings were done in the IC by Dr. Mike Anderson and Prof. Young. Secondly, the lab of Dr. Israel Nelken at Hadassah Medical School of the Hebrew University in Jerusalem, where recordings were conducted in the auditory cortex, the auditory thalamus and the inferior colliculus by Omer Bar-Yosef, Dina Farkas, Liora Las, Nachum Ulanovski and Dr. Nelken.

A detailed description of the experimental methods is given in [Bar-Yosef *et al.*, 2001]. In what follows, a brief review of these is provided.

Animal preparation

Extracellular recordings were made in primary auditory cortex of nine halothane-anesthetized cats, in the medial geniculate body of two halothane- anesthetized cats and inferior colliculus of nine isoflurane-anesthetized and two halothane-anesthetized cats. Anesthesia was induced by ketamine and xylazine and maintained with halothane (0.25-1.5 percent, all cortex and MGB cats, and 2 IC cats) or isoflurane (0.1-2 percent 9 IC cats) in 70 percent N_2O. Breathing rate, quality, and CO_2 levels were continuously monitored. In case of respiratory resistance, the cat was paralyzed with pancuronium bromide (0.05-0.2 mg given every 1-5 hr, as needed) or vecuronium bromide (0.25 mg given every 0.5-2 hr). Cats were Anesthesized using standard protocols authorized by the committee for animal care and ethics of the Hebrew University - Hadassah Medical School (AI, MGB and IC recordings) and Johns Hopkins University (IC recordings).

Electrophysiological recordings

Single neurons were recorded using one to four glass-insulated tungsten microelectrodes micro-electrodes. Each electrode was independently and remotely manipulated using a hydraulic drive (Kopf) or a four-electrode electric drive (EPS; Alpha-Omega, Nazareth, Israel). The electrical signal was amplified (MCP8000; Alpha-Omega) and filtered between 200 Hz and 10 kHz. The spikes were sorted online using a spike sorter (MSD; Alpha-Omega) or a Schmitt trigger. All neurons were well separated. The system was controlled by a master computer, which determined the stimuli, collected and displayed the data on-line, and wrote the data to files for off-line analysis. MGB neurons were further sorted off line

Most of the neurons whose analysis is described below were not recorded simultanousely.

Acoustic stimulation

The cat was placed in a soundproof room (Industrial Acoustics Company 1202). Artificial stimuli were generated digitally at a rate of 120 kHz, converted to analog voltage (DA3-4; Tucker-Davis Technologies), attenuated

(PA4; Tucker-Davis Technologies), and electronically switched with a linear ramp (SW2; Tucker-Davis Technologies). Natural stimuli and their modifications were prepared as digital sound files and presented in the same way, except that the sampling rate was 44.1 kHz. Stimuli were delivered through a sealed calibrated acoustic system (Sokolich) to the tympanic membrane. Calibration was performed in situ by probe microphones (Knowles) precalibrated relative to a Brüel and Kjær microphone. The system had a flat ($\pm$10 dB) response between 100 Hz and 30 kHz. In the relevant frequency range for this experiment (2-7 kHz), the system was even flatter (the response varied by less than $\pm$5 dB in all but one experiment, in which the variation was $\pm$8 dB). These changes consisted of relatively slow fluctuations as function of frequency, without sharp peaks or notches.

Anatomical appraoch

In AI, penetrations were performed over the whole dorso-ventral extent of the appropriate frequency slab (between about 2 and 8 kHz). In MGB, all penetrations were vertical, traversing a number of iso-frequency laminae, and most recording locations were localized in the ventral division. In IC vertical penetrations were used in all experiments except one, in which electrode penetrations were performed at a shallow angle through the cerebellum, traversing the nucleus in a caudo-rostral axis. We tried to map the full medio-lateral extent of the nucleus, but in each animal only a small number of electrode penetrations were performed. Based on the sequence of best frequencies along the track, the IC recordings are most likely in the central nucleus.

1.4 Summary of our approach

With over 100 years of neuroscience research using electrophysiological experiments, how can we hope to innovate, and gain a deeper understanding of the sensory systems?

Our approach is based on combining several ingredients. First, we use natural and complex stimuli, reflecting our belief that interesting properties of high level processing (presumably taking place in the auditory cortex) can be revealed in the responses to such stimuli. Such properties however cannot be discovered using standard linear methods.

Secondly, electrophysiological recordings from a series of auditory processing stations allows us to compare the representations of these complex stimuli, and the way they change along the processing hierarchy, thus reflecting the computational processes that the system applies. Our goal is to identify design principles that underlie the changes in these representations.

Thirdly, we use information theoretic measures to quantify how auditory cells interact to represent stimuli, and develop information theoretic methods to study what the cells represent.

Our belief is that the combination of these components can reveal novel evidence about the principles that underly auditory neural coding.

Bibliography

[Aertsen and Johannesma, 1981] A.M.H.J Aertsen and P.I.M Johannesma. The spectro-temporal receptive field: A functional characteristics of auditory neurons. *Biological Cybernetics*, 42:133–143, 1981.

[Atick and Redlich, 1990] J.J. Atick and A.N. Redlich. Towards a theory of early visual processing. *Neural Computation*, 2:308–320, 1990.

[Atick, 1992] J.J. Atick. Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 3:213–251, 1992.

[Attneave, 1954] F. Attneave. Some information aspects of visual perception. *Psychological Review*, 61:183–193, 1954.

[Bar-Yosef *et al.*, 2001] O. Bar-Yosef, Y. Rotman, and I. Nelken. Reponses of neurons in cat primary auditory cortex to bird chirps: Effects of temporal and spectral context. *Journal of Neuroscience*, 2001.

[Barlow, 1959a] H.B. Barlow. The coding of sensory messages, chapter XIII. In Thorpe and Zangwill, editors, *Current Problems in Animal Behavior*, pages 330–360. Cambridge University Press, 1959.

[Barlow, 1959b] H.B. Barlow. Sensory mechanisms, the reduction of redundancy, and intelligence. In *Mechanisation of thought processes*, pages 535–539. Her Majesty's stationary office, London, 1959.

[Barlow, 1961] H.B. Barlow. Possible principles underlying the transformation of sensory messages. In W.A. Rosenblith, editor, *Sensory communication.* MIT Press, Cambridge, MA, 1961.

[Barlow, 2001] H.B. Barlow. Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12:241–253, 2001.

[Becker, 1996] S. Becker. Mutual information maximization: Models of cortical self organization. *Network: Computation in Neural Systems*, 7:7–31, 1996.

[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.

[Chechik, 2003] G. Chechik. Spike-timing dependent plasticity and relevant mutual information maximization. *Neural Computation*, 15(7):1481–1510,, 2003.

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

[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.

[Eggermont *et al.*, 1983] J.J. Eggermont, P.I.M Johannesma, and A.M.H.J Aertsen. Reverse correlation methods in auditory research. *Quarterly Reviews of Biophysics*, 16:341–414, 1983.

[Ehret and Merzenich, 1988] G. Ehret and MM. Merzenich. Complex sound analysis (frequency resolution filtering and spectral integration) by single units of the inferior colliculus of the cat. *Brain Research*, 13(2):139–163, 1988.

[Fishbach *et al.*, 2001] A. Fishbach, I. Nelken, and Y. Yeshurun. Auditory edge detection: A neural model for physiological and psychoacoustical responses to amplitude transients. *The Journal of Neurophysiology*, 85:2303–2323, 2001.

[Haykin, 1999] S. Haykin. *Neural Networks: A Comprehensive Foundation.* Prentice Hall, Upper Saddle River, New Jersey, 2nd edition, 1999.

[Heil and Irvine, 1996] P. Heil and D.R.F Irvine. On determinants of first-spike latency in auditory cortex. *NeuroReport*, 7:3073–3076, 1996.

[Heil, 1997] P. Heil. Auditory cortical onset responses revisited I. first-spike timing. *The Journal of Neurophysiology*, 77(5):2616–41, May 1997.

[Linsker, 1988] R. Linsker. Self organization in a perceptual network. *Computer*, 21:105–117, 1988.

[Linsker, 1989] R. Linsker. An application of the principle of maximum information preservation to linear systems. In D.S. Touretzki, editor, *Advances in Neural Information Processing Systems*, volume 1, pages 186–194, San Mateo, CA, 1989. Morgan Kaufman.

[Linsker, 1992] R. Linsker. Local synaptic learning rules suffice to maximize mutual information in a linear network. *Neural Computation*, 4:691–702, 1992.

[Linsker, 1997] R. Linsker. A local learning rule that enables information maximization for arbitrary input distribution. *Neural Computation*, 9:1661–1665, 1997.

[Middlebrooks *et al.*, 2002] J. Middlebrooks, L. Xu, S. Furukawa, and B.J. Mickey. Location signaling by cortical neurons. In D. Oertel, R.R. Fay, and A. N. Popper, editors, *Integrative Functions in the Mammalian Auditory Pathway*, volume 15 of *Handbook in Auditory Research*, pages 319–357. Springer-Verlag, New York, 2002.

[Miller, 1956] G.A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1956.

[Nadal and Parga, 1994] J.P. Nadal and N. Parga. Nonlinear neuron in the low noise limit: A factorial code maximizes information transfer. *Network: Computation in neural systems*, 5:565–581, 1994.

[Nadal *et al.*, 1998] J.P. Nadal, N. Brunel, and N. Parga. Nonlinear feed-forward networks with stochastic outputs: Infomax implies redundancy reduction. *Network: Computation in neural systems*, 9:207–217, 1998.

[Nelken, 2002] I. Nelken. Feature detection by the auditory cortex. 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*. Springer-Verlag, 2002.

[Palombi and Caspary, 1996] P.S. Palombi and D.M. Caspary. GABA inputs control discharge rate primarily within frequency receptive fields of inferior colliculus neurons. *The Journal of Neurophysiology*, 75(6):2211–2219, 1996.

[Popper and Fay, 1992] A.N. Popper and R.R. Fay. *The mammalian auditory pathway: Neurophysiology*. Springer Verlag, New York, 1992.

[Schreiner and Langner, 1997] C.E. Schreiner and G. Langner. Laminar fine structure of frequency organization in the auditory midbrain. *Nature*, 388:383–386, 1997.

[Spangler and Warr, 1991] K.M. Spangler and W.B. Warr. The descending auditory systems. In R.A. Altschuler, R.P. Bobbin, B.M. Clopton, and D.W. Hofmann, editors, *Neurobiology of Hearing: The central Auditory System*, pages 27–45. Raven, New York, 1991.

[Ukrainec and Haykin, 1996] A.M. Ukrainec and S. Haykin. A modular neural network for enhancement of cross-polar radar targets. *Neural Networks*, 9:143–168, 1996.

[Uttley, 1970] A.M. Uttley. *Information transmission in the nervous system*. Academic press, London, 1970.