Automated detection of mouse scratching through audio recording

Peter Elliott (Max G'Sell, Valerie Ventura)

Draft/Outline

1 Introduction

`This is not a true introduction but may be a helpful reference`
The paper I will be writing will be about methodology for automated detection of scratching by caged mice. I have worked on a method to take an audio recording of a mouse and return labeled time points when scratching occurs. The paper will need to explain why this work matters: why record scratches and why bother with automation? The paper will also need to explain how my method works and how well it works. Since the work is primarily of interest to non-statisticians, many of the techniques I've used will likely be unfamiliar to much of the audience.

The ultimate goal of the project is to aid efforts to research itch and pain. These sensations are often studied through behavioral observation. In the case of itch, mice may be injected with a chemical agent that induces itch then scratching behavior is observed over time. Quantity of scratching is used to assess how much itch the mouse felt (unfortunately, they can't tell us directly). This can be used to measure the effect of various interventions, such as genetic modifications to the mice or skin treatment. This observation helps us understand the mechanisms by which itch and pain work. Improved understanding and the ability to test interventions may eventually lead to improved treatment options for humans.

Currently, observations of mice are typically done manually. That is, someone has to sit, watch and count. As a result, the sample sizes collected are often very small. Furthermore, often only a small fraction of the mouse's response to induced itching is observed - the effect of the chemical agent continues long after observation stops. These problems would be fixed by automated observation. This would allow many mice to be observed simultaneously and would allow for extended observation times with negligible extra expense to the scientist's time.

The method the paper will propose uses audio data to automate the observation process. The method requires that a single mouse be placed in a sound-proofed cage with a microphone for the duration of the experiment. The resulting audio

data is analyzed using a two-stage process: first segmentation into candidate time intervals then classification. The segmentation process makes use of scratching's characteristic rhythmic pattern. Time intervals with regular peaks in energy are marked as candidates. The classification is done using a random forest trained using features we have designed. These features attempt to encode both time and frequency (i.e. audio pitch) information. We will discuss the design of these features.

The method was developed using data collected by the Ross Neuroscience Lab at the University of Pittsburgh, and we use that data to assess the accuracy of our method. The classifier is not perfect, so the statistical implications of mislabeled scratches will need to be discussed.

2 Methodology

Mouse scratching occurs in small sets of rapid swipes, which we refer to as *scratch bouts*. The goal of the procedure is to detect when these scratch bouts occur.

The proposed procedure for scratch bout detection has three main steps. First, data is collected as an audio recording of a single mouse in its cage. Second, the recording is segmented into time intervals that are candidates to be labeled as scratches. Finally, a classifier is used to label the candidate time intervals.

2.1 Data collection

Put a mouse in a cage with a microphone. `there's more to say here`

2.2 Segmentation

After data is collected, the recording is broken up into time intervals ~~in which a scratch bout may have occurred.~~ At this stage of processing, we wish to capture the vast majority of scratch bouts while minimizing the number of scratch-free periods that will be considered for labeling.

Our strategy for segmentation relies on the characteristic rhythm caused by the series of swipes that occur within a scratch bout. Within a scratch bout, swipes tend to occur with a period of approximately 50ms. Exploratory analysis of the data also reveals that sound caused each swipe occupies a broad frequency band. Additionally, while lower frequencies are often contaminated by other sources of sound, higher frequencies tend to uncontaminated. We therefore wish to locate time points with strong power in higher frequencies and check that these peaks in power occur close to 50ms apart.

Segmentation goes as follows:

1. Create a spectrogram from the audio recording. We use time bins of 128 samples (approximately 3ms) with an overlap of 96 samples and a Gaussian windowing function (`reference may be needed here`).

2. For each time bin, sum the modulus of Fourier coefficients corresponding to frequencies over 10kHz. This measures the power in the frequency range above 10kHz.

3. Apply a triangular kernel smoother to the time series of power measurements. The kernel shape matches the expected shape in the time series of peaks caused by scratches. We use a bandwidth that corresponds to approximately 8ms.

4. Find time points in the smoothed time series that (a) have maximal power within 25ms, and (b) have power greater than the minimal time bin within 25ms by some chosen parameter h.

5. Group sequences of time points found in part 4 that have between-adjacent-point distances under 120ms. Groups that include at least three time points are marked as candidate time intervals.

Most of the parameters used are a reflection of exploratory analysis. The 10kHz cutoff is where sound contamination starts to be rare. The bandwidth of the kernel smoother is chosen to mirror the typical duration of a swipe. The 25ms range in which a time point must be a maximum is a reflection of the minimum amount of time there can be between swipes in a scratch bout. The 120ms distance is used to avoid splitting a scratch bout in the event that one of its swipes is not detected.

Sensitivity of the segmentation process is controlled by the choice of h. Higher values of h require that time bins be more pronounced maxima to be considered as candidate swipes, and therefore fewer candidate time intervals will be chosen. An appropriate value for h can be chosen using grid search with a training set.

2.3   Classification

Labeling the candidate time intervals requires the design of features that capture important characteristics of the data as well as the selection of a classification method capable of distinguishing between the class-conditional distributions. Many of the features we use try to find and describe rhythmic patterns in the data. We also make use of audio pitch (frequency) information via the Fourier transform. For classification, we use a Random Forest, which is capable of using many potentially correlated features.

2.3.1   Feature design

Recall that for the purpose of segmentation, the waveform representation of the audio was transformed via summation of Fourier coefficients and kernel smoothing to give a new time series representation that contained useful information. There are other useful transformations that can also be made. We can extract common features from these transformations to get a large body of potentially informative features.

3

Transformations

We use two main types of transformations. The first involves convolution of power measurements across time with a template. The transformation used for segmentation (described in steps 1-3 in that section `should be referenced more cleanly`) is an example - in that case the triangular kernel acts as the template. The second type of transformation involves computing a spectrogram for the time interval and calculating a correlation or inner product with a periodogram template at each time bin. The transformations we use are as follows:

1. Convolve the power above 10kHz with a triangular kernel.

2. Convolve the power above 15kHz with a triangular kernel.

3. Convolve the power above 20kHz with a triangular kernel.

4. Calculate the correlation of the spectrum at each time point with an average of periodograms drawn from a sample of scratches.

5. Calculate the inner product of the spectrum at each time point with an average of periodograms drawn from a sample of scratches.

`I'm using spectrum and periodogram sort of interchangeably which is probably bad`
  `There are a couple other transformations that need to be described`

Common features

For each transformation, we wish to extract a set of features. For scratch bouts, we expect to again see a rhythmic pattern of peaks in the transformed time series. The features we extract should describe the set of peaks found in the candidate time interval. Peaks in the transformed time series can be found following step 4 in the segmentation process. We then extract the following features:

1. The number of peaks

2. The average time between peaks

3. The standard deviation of times between peaks

4. Summary statistics (mean, median, minimum, maximum, and standard deviation) of the powers at each peak

5. Summary statistics of the full width at half maximum of each peak

6. Summary statistics of the transformed time series itself

In some cases, the feature to extract is not defined. We cannot calculate the standard deviation of the times between peaks, for example, if the transformed time series only has two peaks. In these cases, the feature can be coded as -1. `There might be a better place to mention this`

Miscellaneous features

- `time`

- `large-binned spectrograms and frequency distributions`

2.3.2 Random Forests

`Needs thought about appropriate level of discussion`
`   benefits to cover`

- `lots of features are thrown in`

- `lots of features are heavily correlated`

- `classes are imbalanced`

- `decision boundary probably nonlinear`

- `vote implementation (R vs.  sklearn)`

3 Results

- `accuracy measures (sensitivity, FDR, ROC?)`

- `analysis of where in the process errors occur (segmentation vs classification)`

- `comparison to previous methods`

- `stereo recording vs mono?`

- `comparison of classifier choices?  (RF, logistic lasso)`

- `should evaluate robustness to parameter choice`

4 Discussion