


Identification of Fixations in Noisy Eye Movement Data via Recursive Subdivision



Jeffrey B. Mulligan¹ and Donald J. Kalar²

¹NASA Ames Research Center, ²San José State University




The Problem

The noise level in a modern video-based eye-tracker is usually larger than the size of the smallest saccades. Thus, it can be challenging to discriminate small eye movements from system noise. Various methods have been proposed, but most require hand-tuning of parameters to achieve good performance. We would like a method that automatically adapts to the statistics of the input signal.

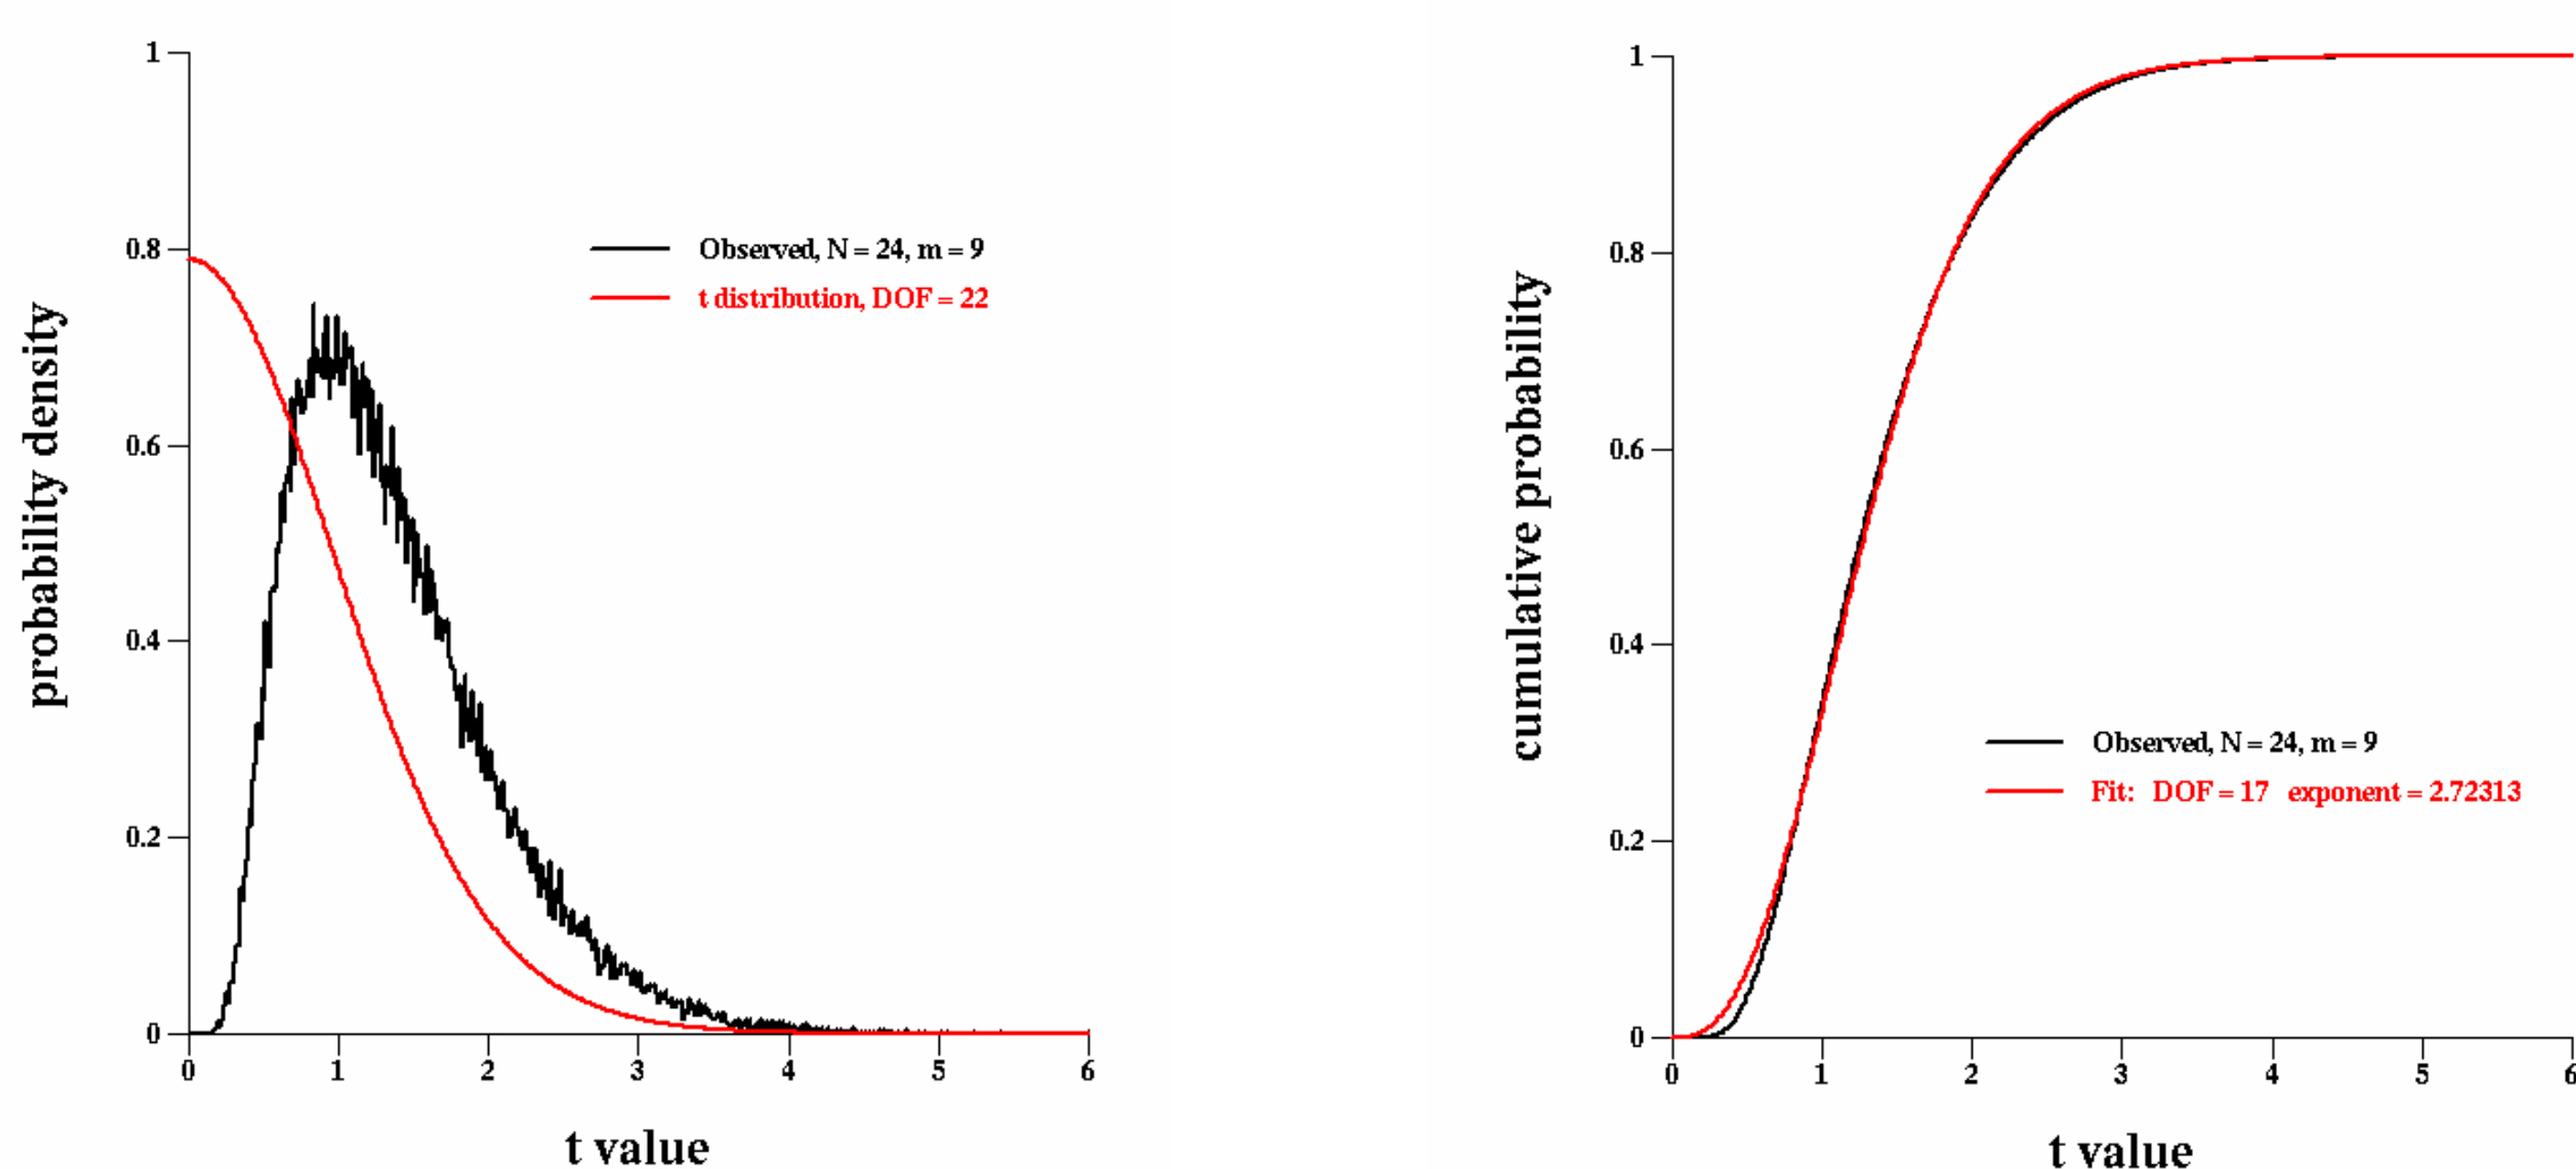
The Proposed Solution

We attempt to fit a signal with a piecewise-constant function. We begin by approximating the signal with just two constant intervals, linked by a saccade. We perform exhaustive search over all possible split points, with the minimum fixation duration a parameter (set to 150 ms.). We approximate each interval with the mean of the samples in that interval, and calculate the root-mean-square (RMS) error. We choose the split point that produces the lowest RMS error, and then apply the procedure recursively to subintervals that are longer than twice the minimum fixation duration.




Statistical testing

After finding the optimal split point, we apply a statistical test to decide whether or not to accept the split. Originally, we used a simple t-test. We soon discovered however, that we were obtaining too many false alarms (when testing with a constant signal corrupted by Gaussian noise). This can be understood when it is noted that we are not performing the t-test at a predefined split point, but rather at the split point that best separates the data. We then tried a bootstrap method, in which we repeated the procedure with many permuted versions of the data, and only accepted the split if the t value was larger than that obtained with the permuted samples sufficiently many times (based on the desired level of significance). Because, the bootstrap method is extremely slow, however, it is desirable to model the distribution of the t-statistic obtained with the procedure.



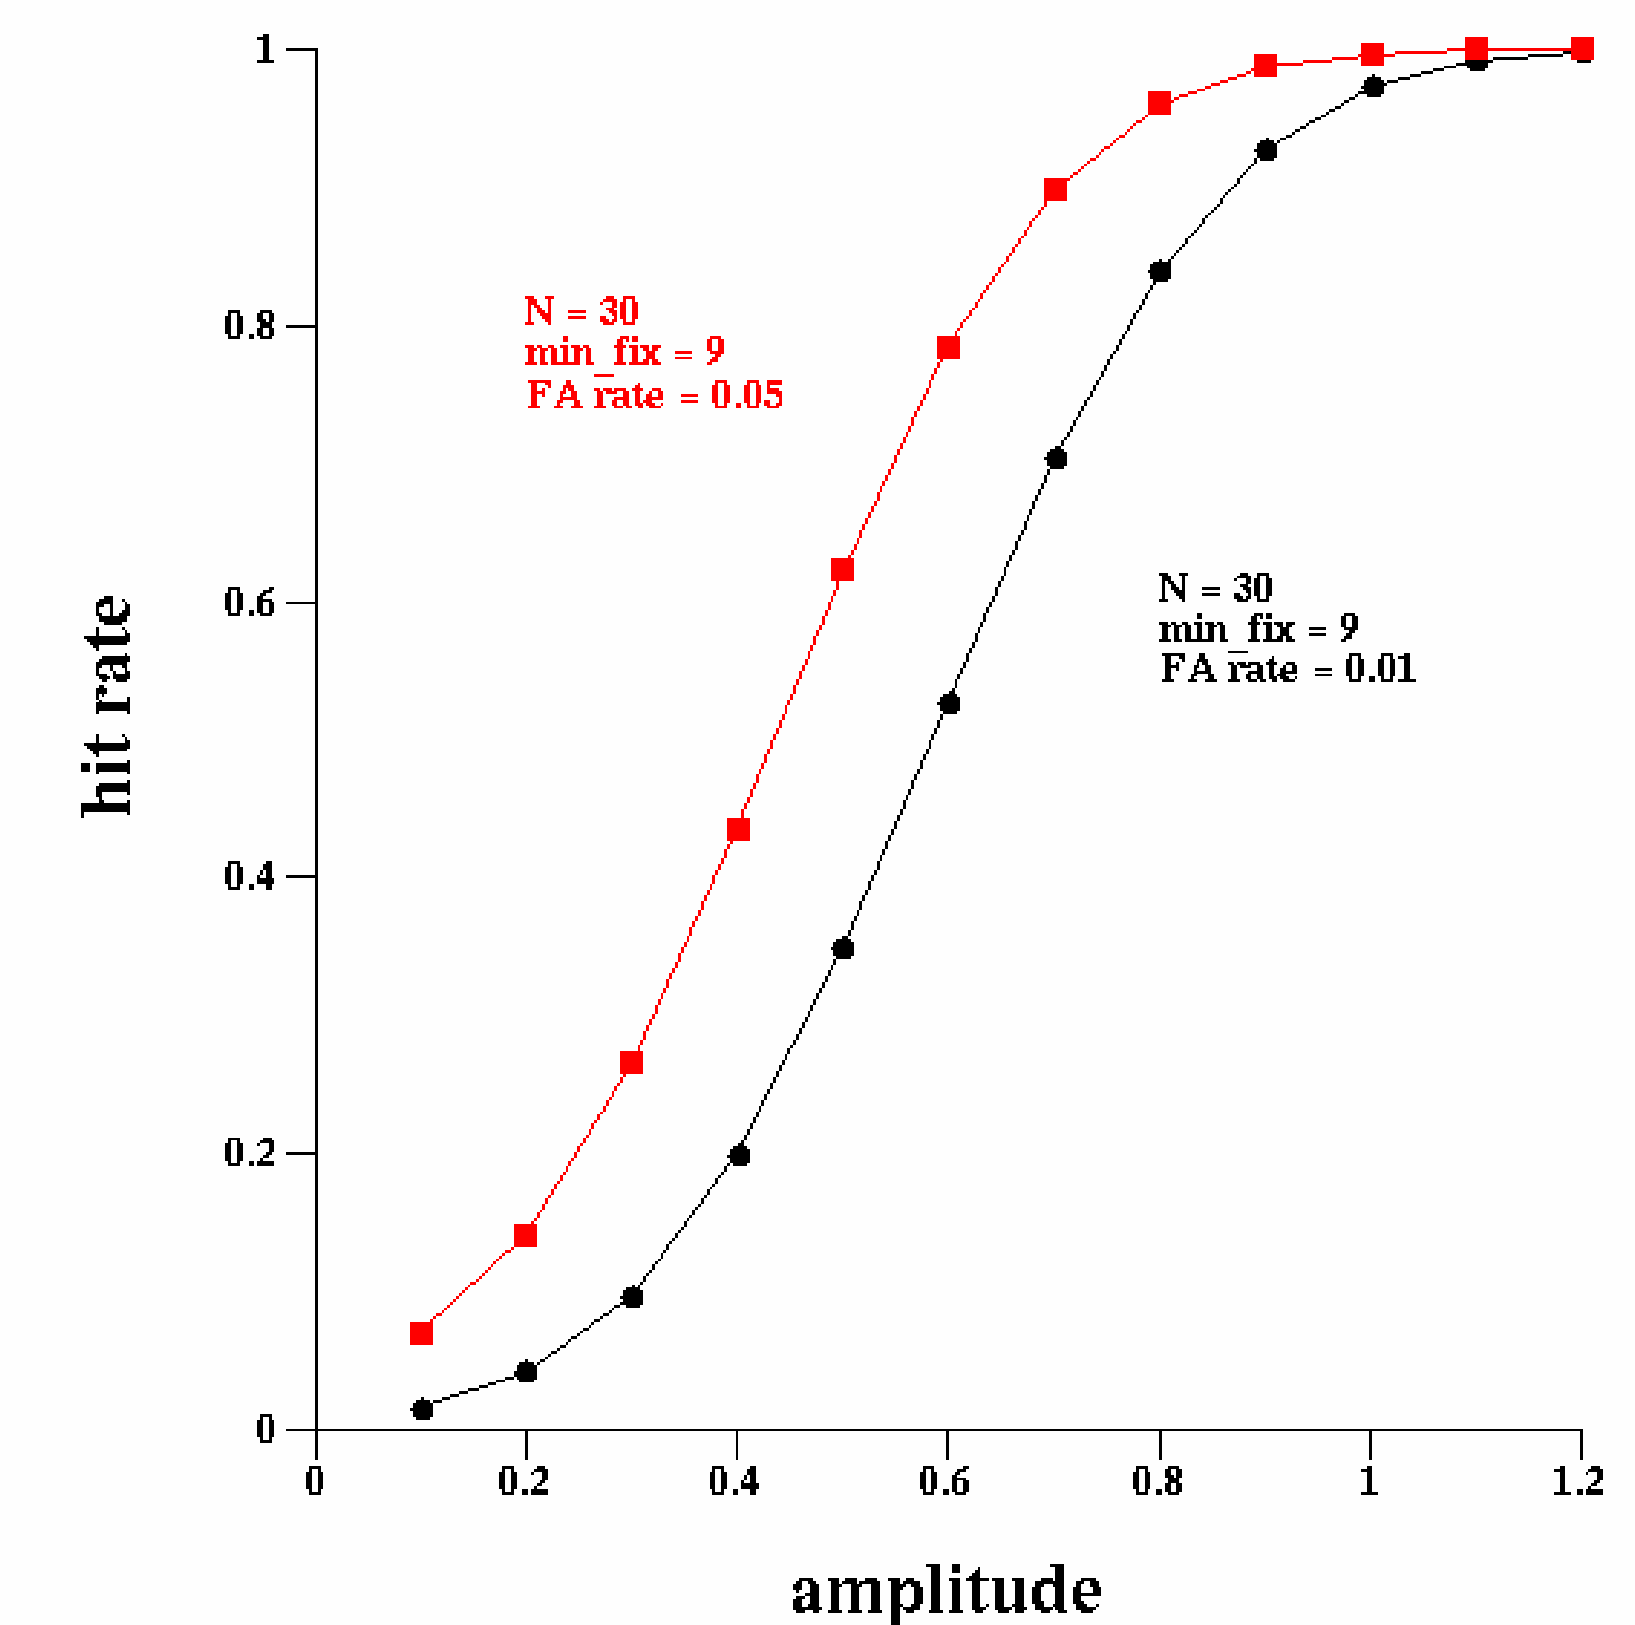
The observed noise distribution is well-fit by the distribution of the maximum of samples from a t distribution, which can be obtained by raising the cumulative t distribution to a power corresponding to the number of samples. In practice, we speed computation by tabulating values of the critical t value as a function of significance level, sequence length, and minimum fixation duration.




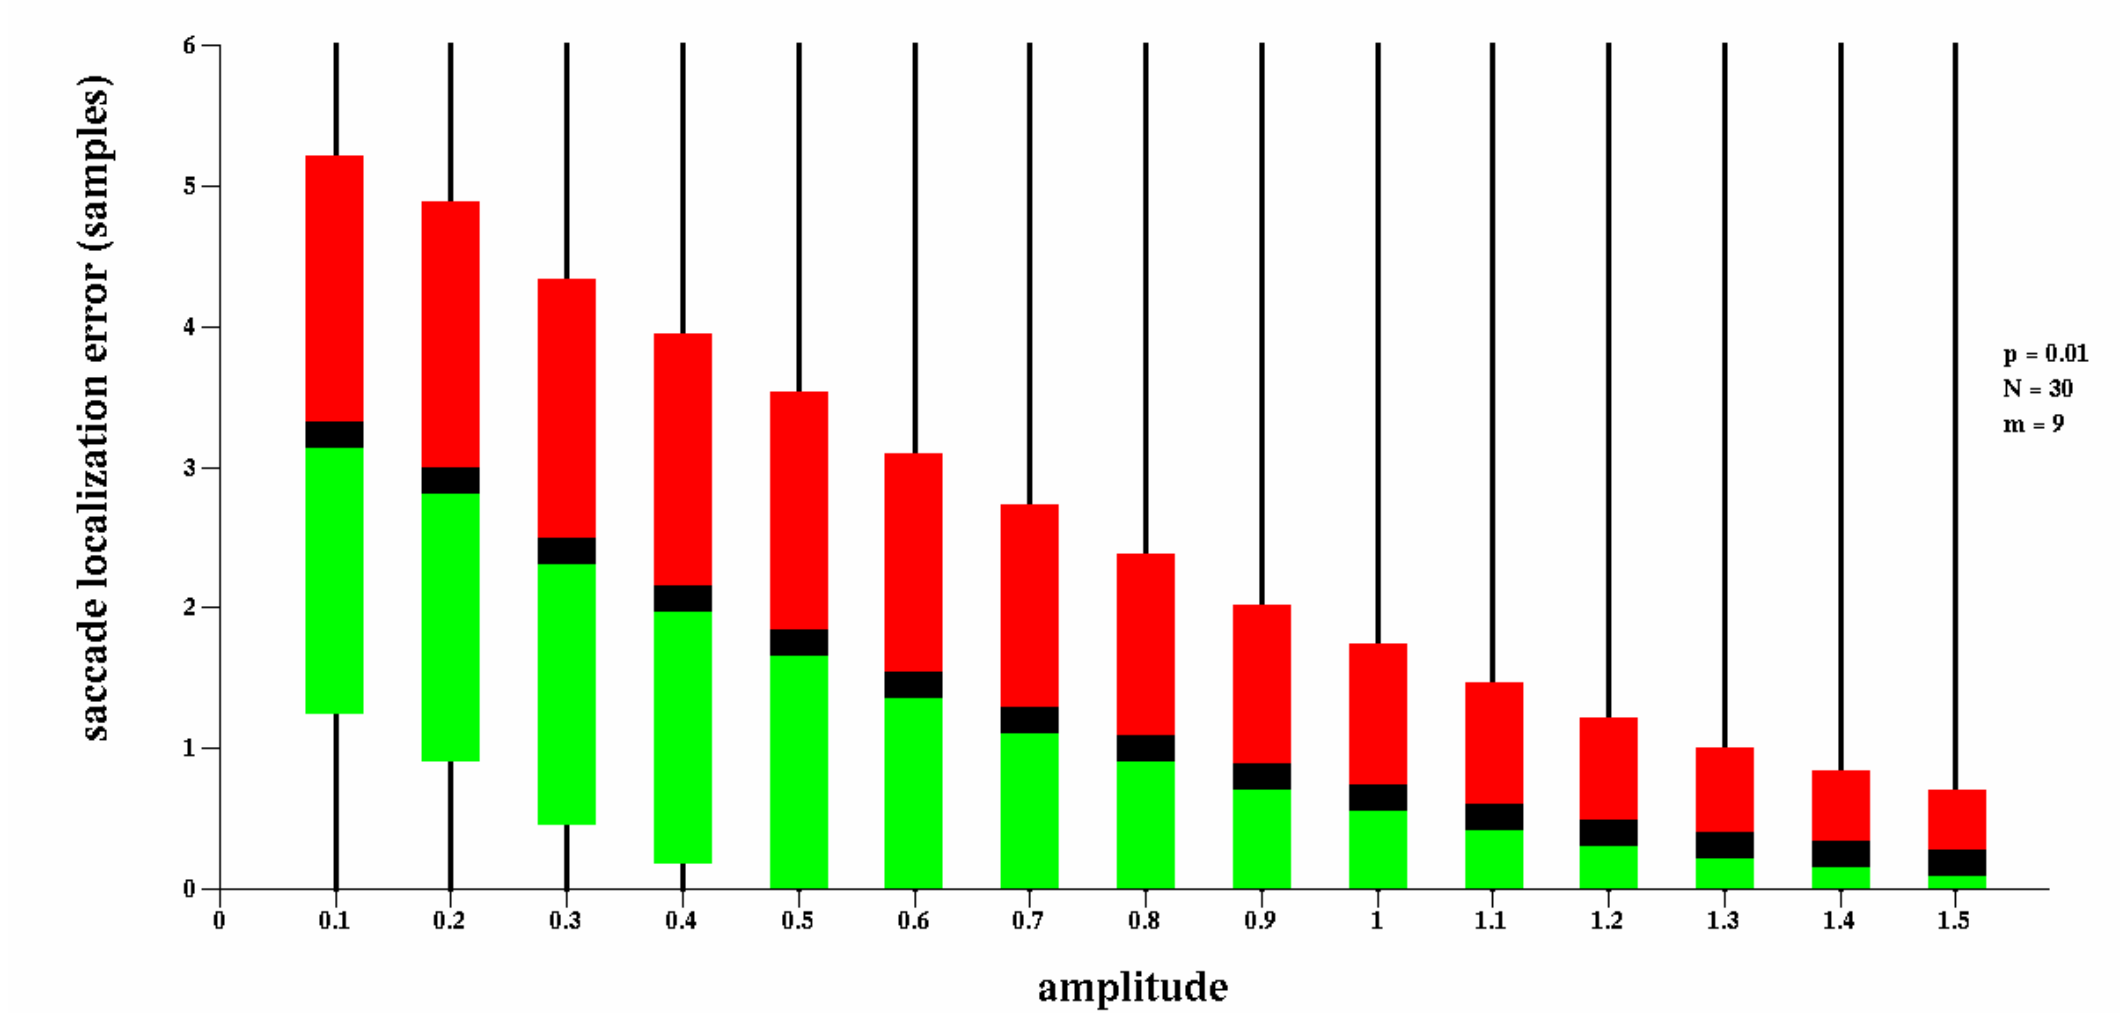

We compute ROC curves from Monte Carlo simulations of known signals corrupted by Gaussian noise with unit variance. In the example shown above, we see that we can detect 62% of the small saccades of an amplitude of half the noise standard deviation if we accept a false alarm rate of 5%. The hit rate drops to 35% for a false alarm rate of 1%.

Sensitivity

The figure to the right shows hit rate as a function of saccade amplitude (in noise units) for two significance levels.





Saccade timing errors



In the analysis presented here, hits include trials where the split was placed at a sample other than the saccade location (according to the ground truth before the addition of noise). Even for easily detectable saccades, the time of occurrence of the inferred saccade can be different from the true time.

Saccade amplitude errors



For the smallest saccades, the estimated amplitude is generally higher than the true amplitude. This is because a certain level difference is required to achieve statistical significance. But sometimes the noise can add to the signal and push it to significance. In the figure to the right, only the left-most points include direction errors (leading to the higher standard deviations, and the downward inflection of the curve).

