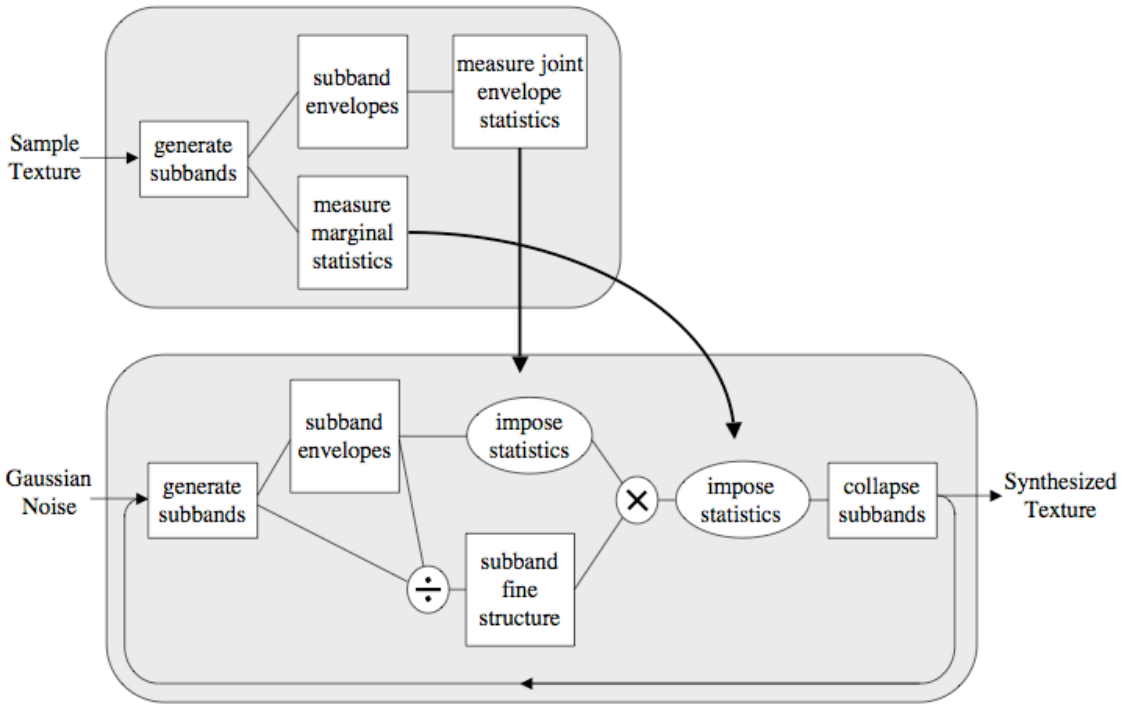

Figure 1. Schematic of texture synthesis algorithm. Because statistical adjustments interfere with each other, and with the process of collapsing and re-generating the subbands, the process is iterated. Over time it generally converges to a stable state, with a signal matching the statistics of the original.

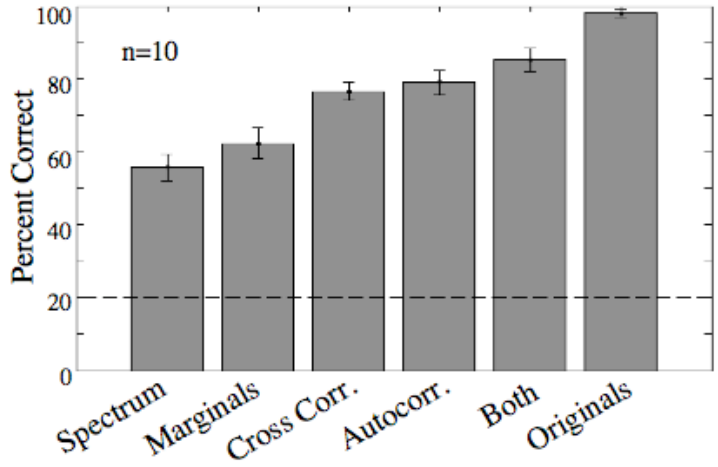

Figure 2. Results of texture identification experiment. Listeners were presented with 5 second sound excerpts, and asked to identify which of 5 categories it belonged to. Each sound was either the original sound, or a version synthesized with various statistical constraints (power spectrum, marginals (variance, skew, kurtosis), cross-band envelope correlation, within band envelope autocorrelation, or both correlations). Plotted data are averaged over 10 subjects. Performance improves markedly as additional statistics are imposed.

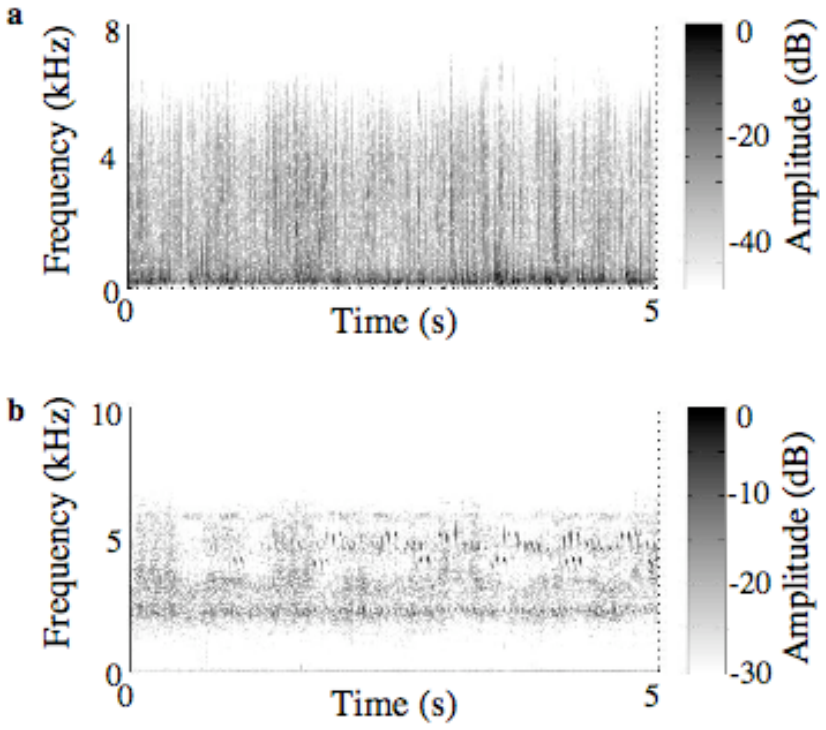

Figure 3. Spectrograms of original recordings of fire (a) and swamp insects (b).

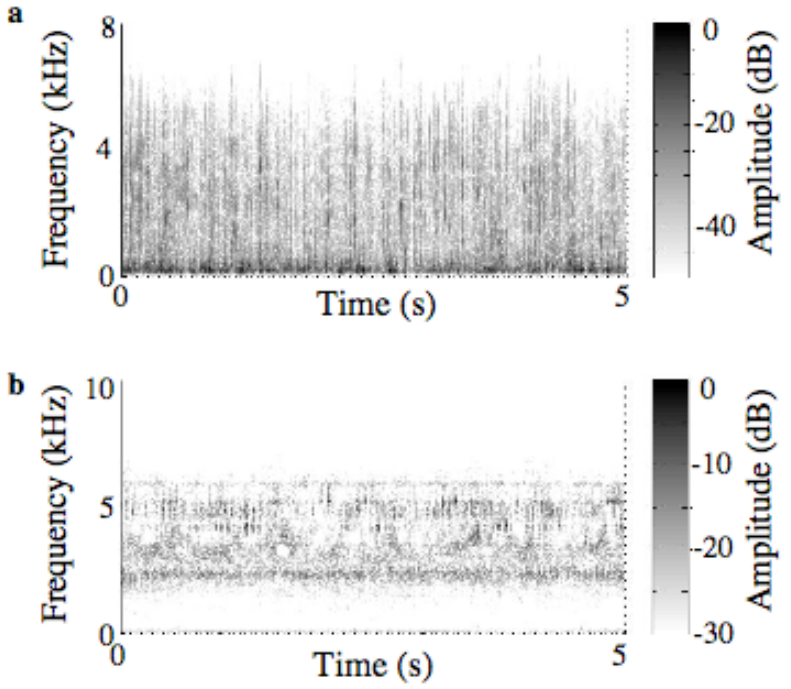

Figure 4. Spectrograms of synthesized versions of fire (a) and swamp insects (b), synthesized using both marginal statistics and envelope correlations.