


A MARKOV-CHAIN MONTE CARLO APPROACH TO MUSICAL AUDIO SEGMENTATION

Christophe Rhodes, Michael Casey, Samer Abdallah, Mark Sandler

Duration Prior




Results

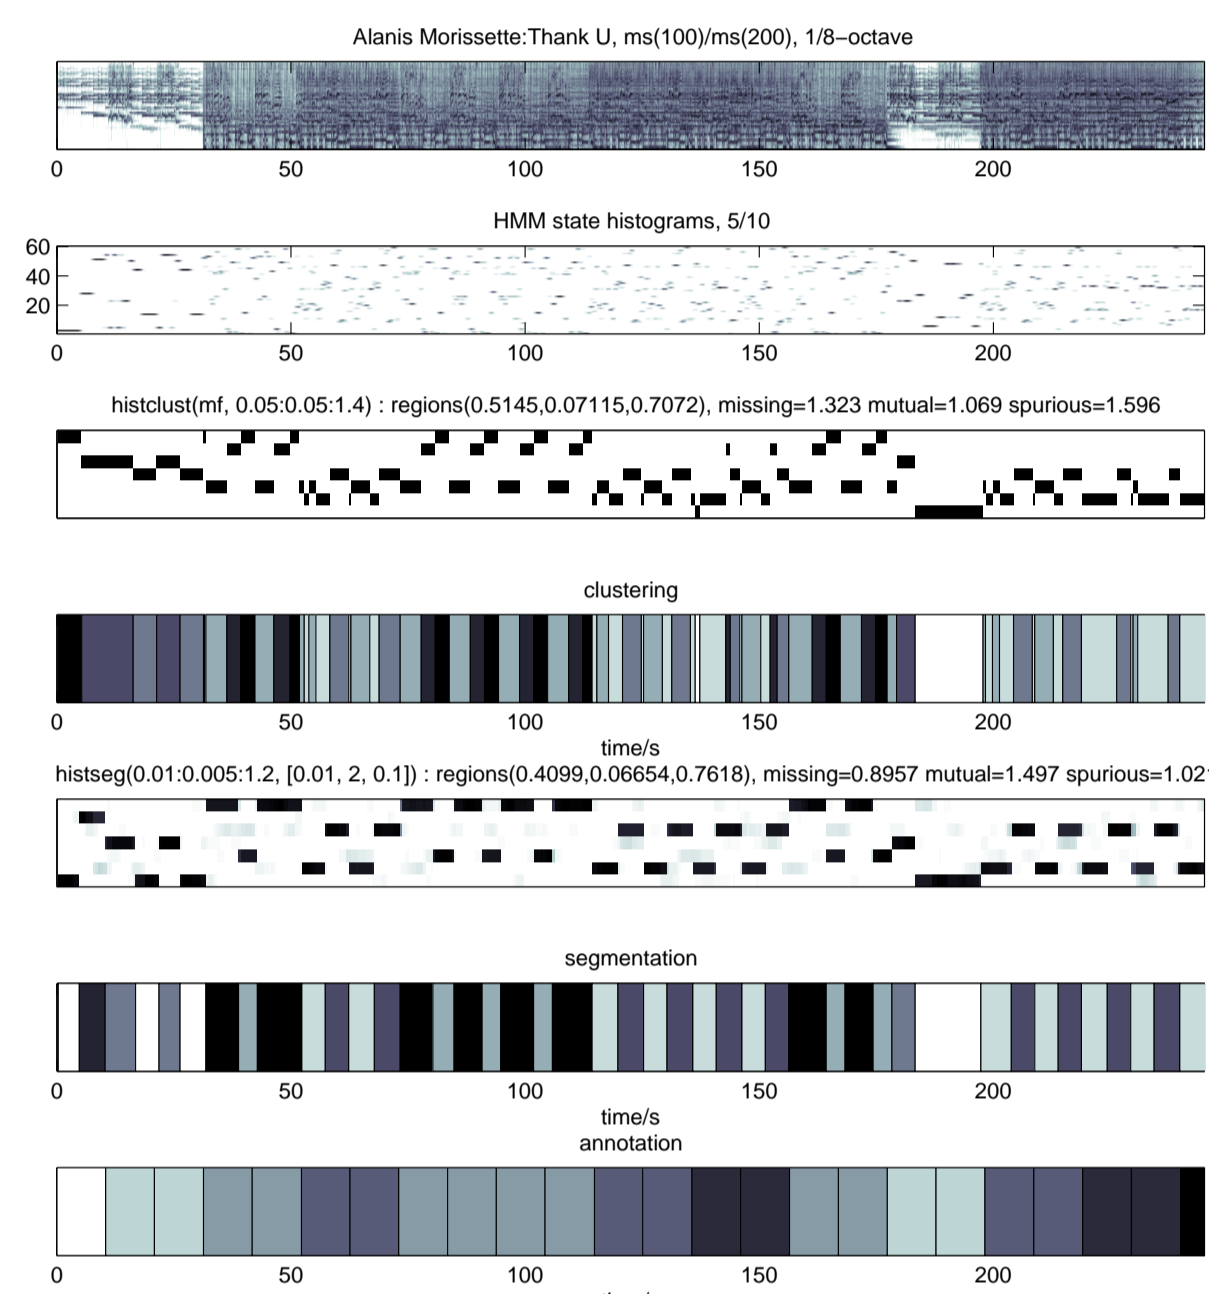
We present some segmentations generated by our segmenter, using our Wolff-Gibbs sampler with a broad prior for segment lengths. The annotations were provided by an expert listener, and are mixtures of functional and audible segmentations.

We used the prior, left, to model our expectations for segment size; note the strong cutoff at short times, and the broad tail for longer times. We plot the F summary statistic for segmentations generated for the same number of classes as labels in the expert annotation.

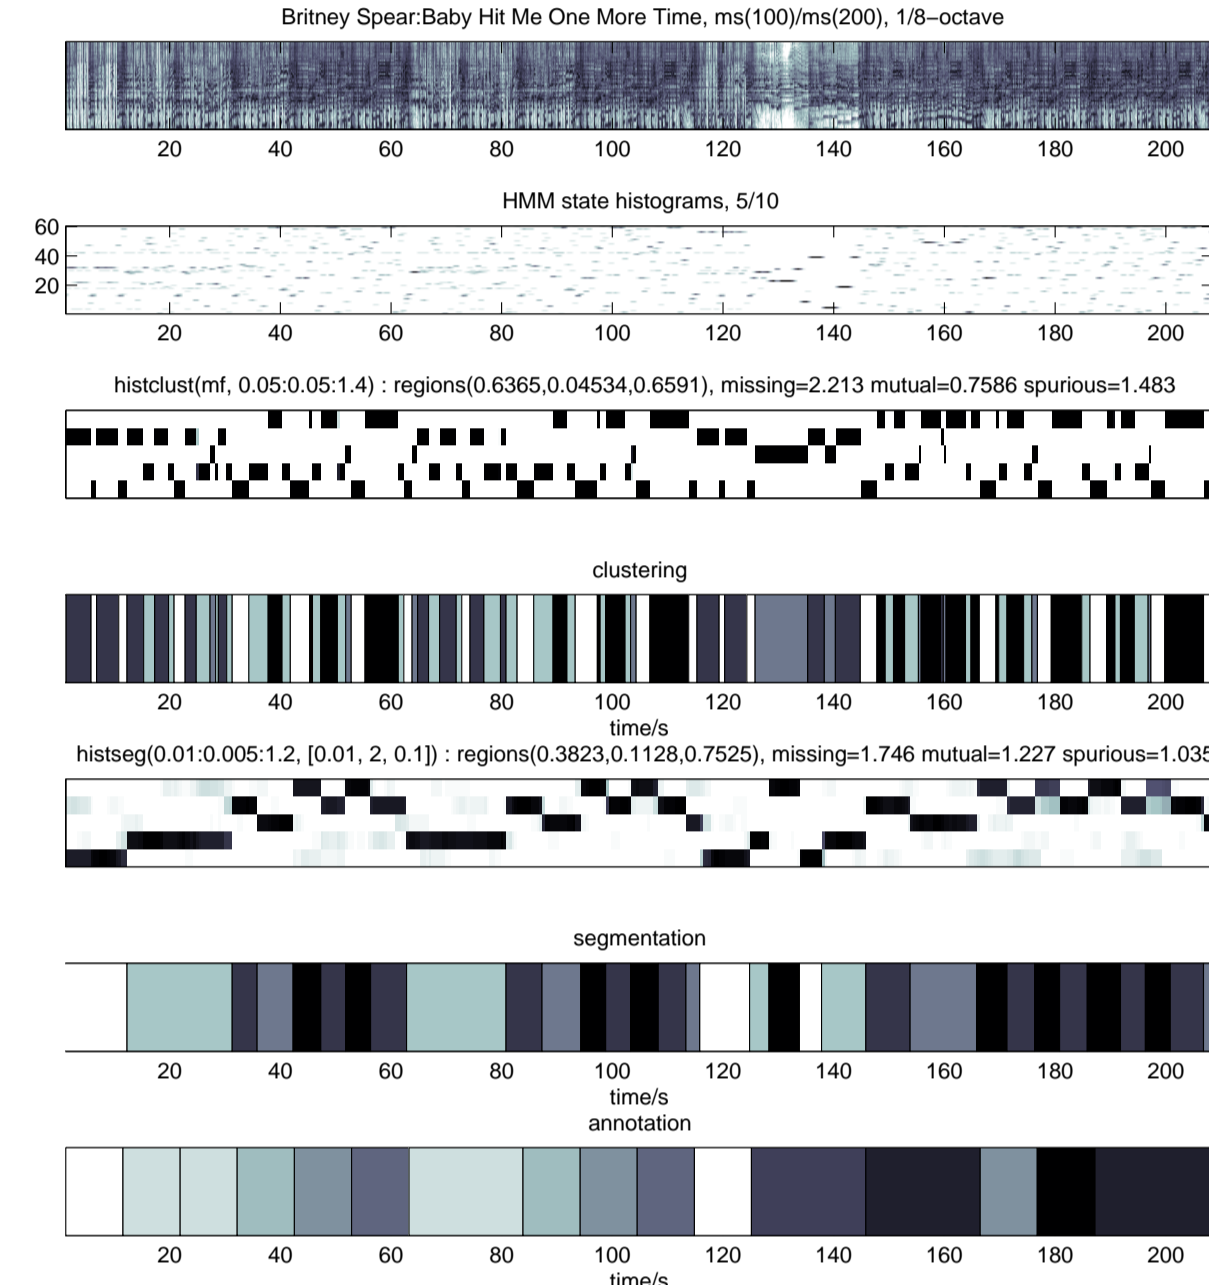
Summary



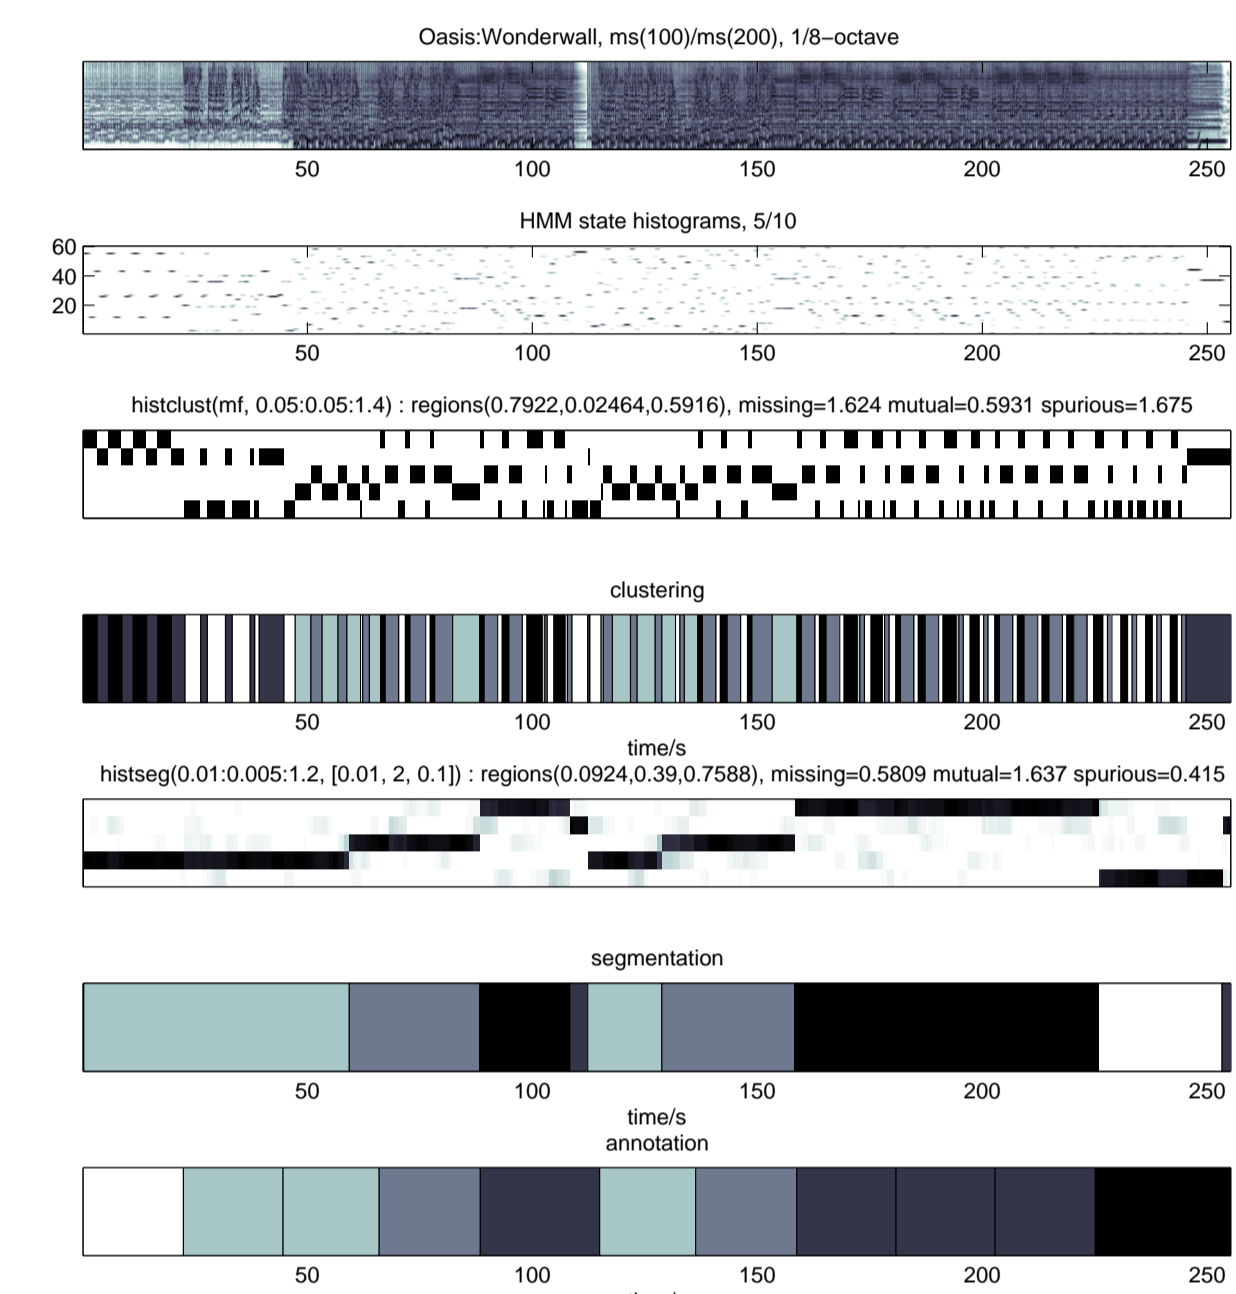
Alanis Morissette: Thank U



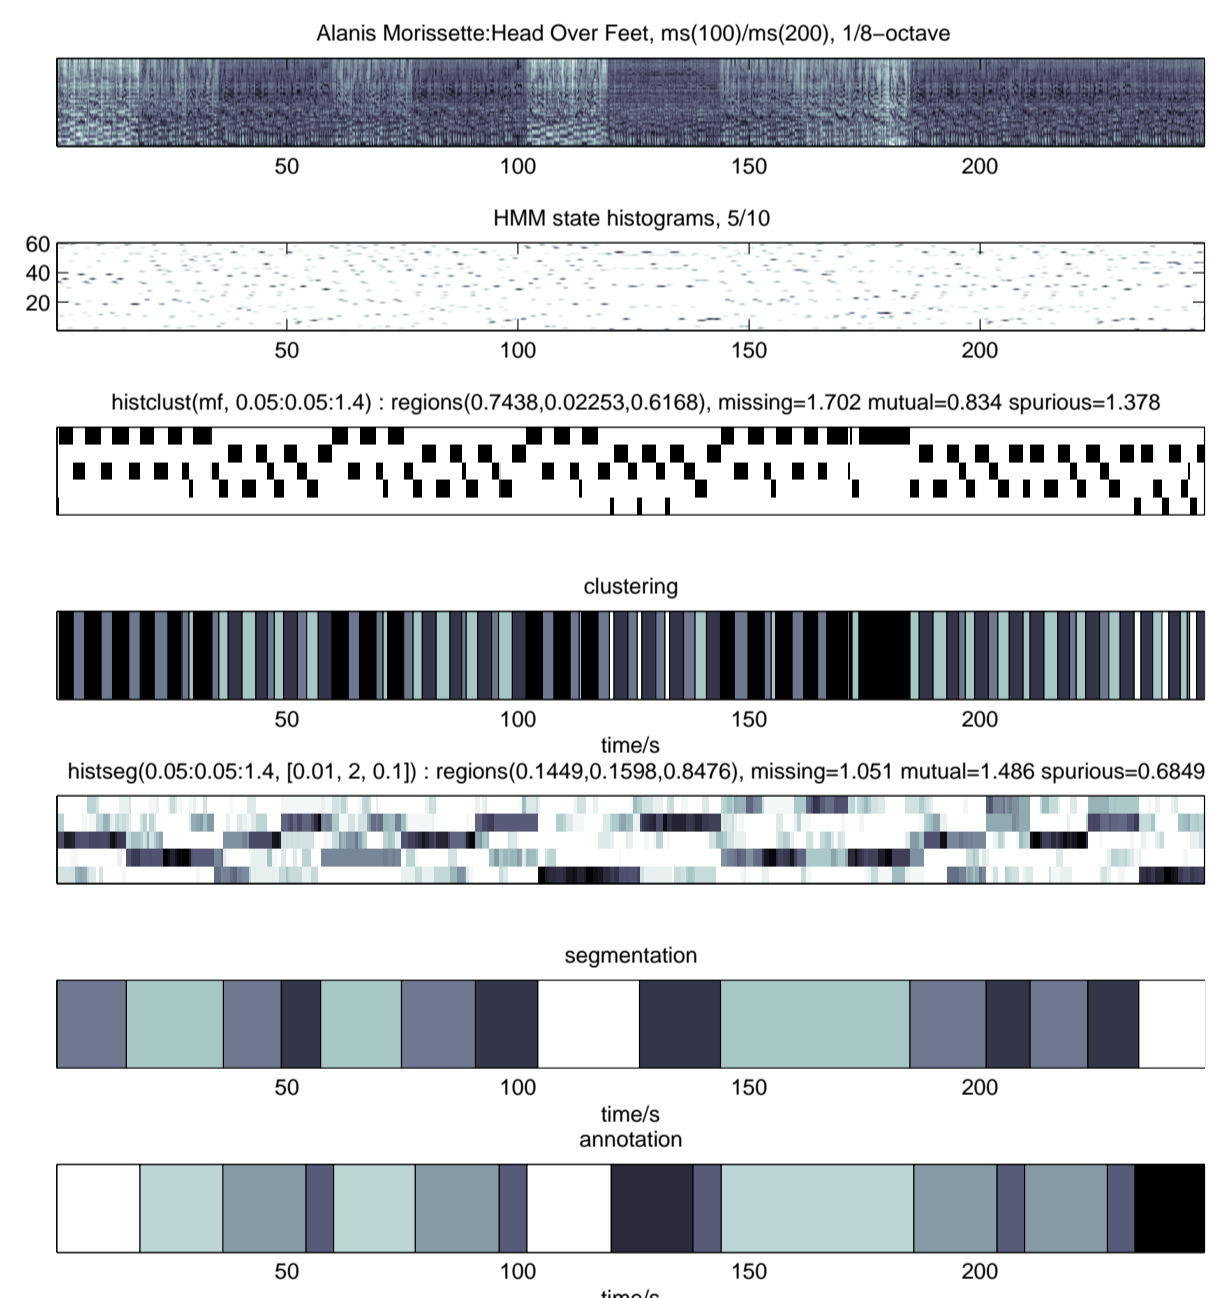
Britney Spears: Hit Me Baby



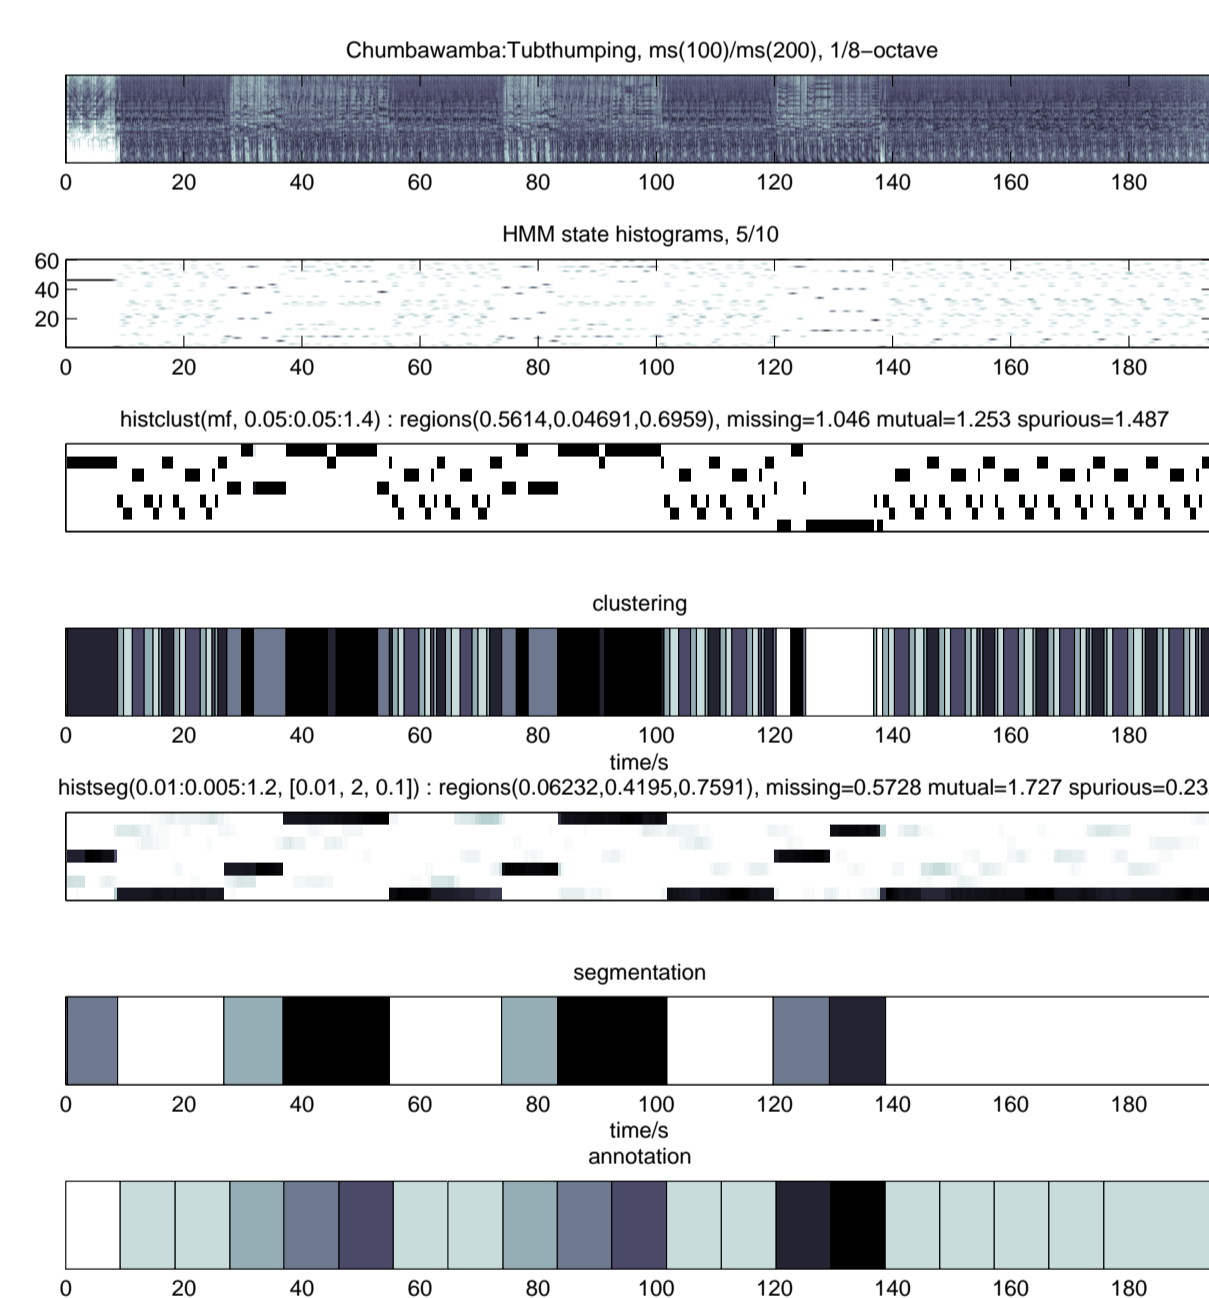
Oasis: Wonderwall



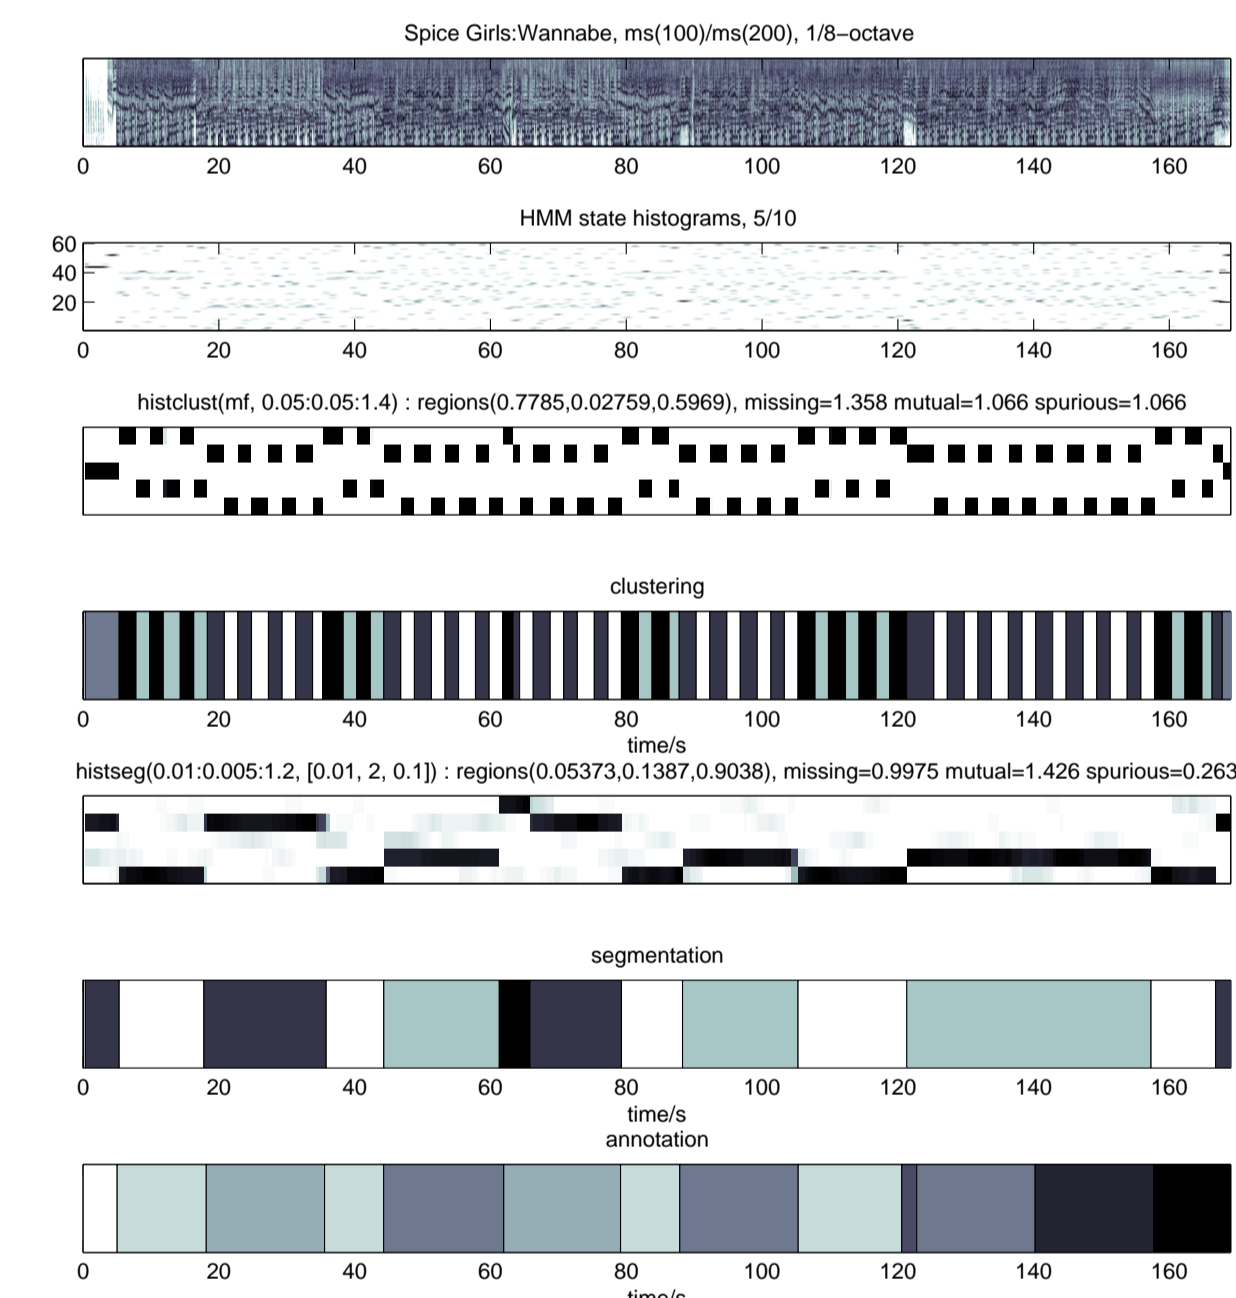
Alanis Morissette: Head Over Feet



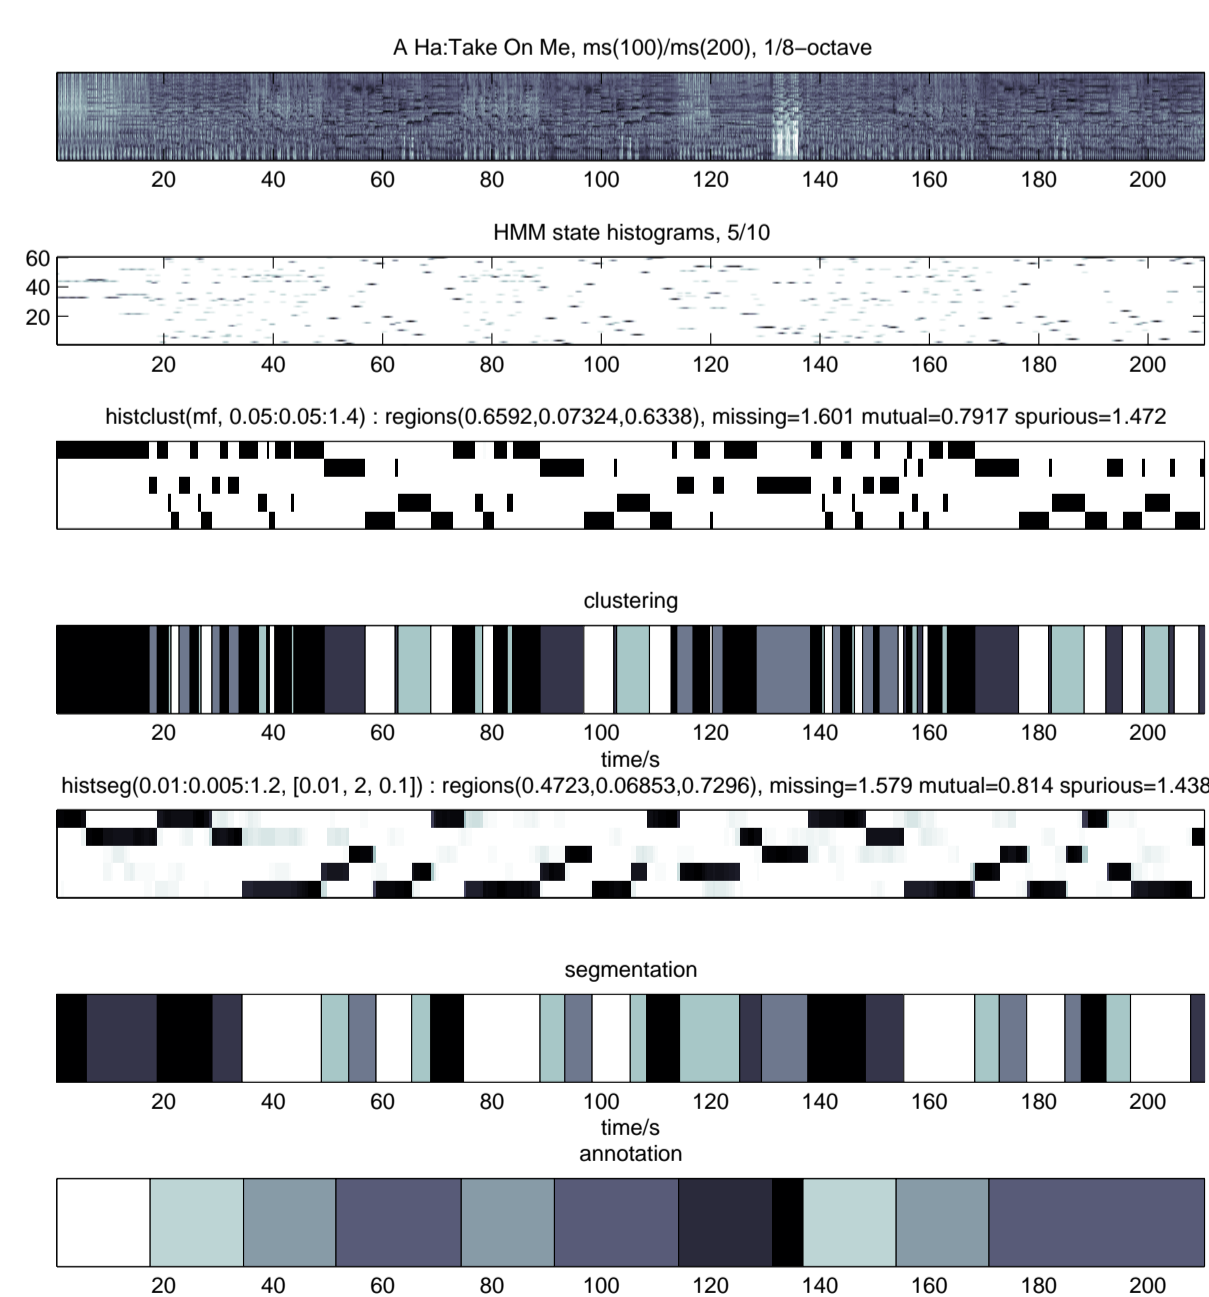
Chumbawumba: Tubthumping



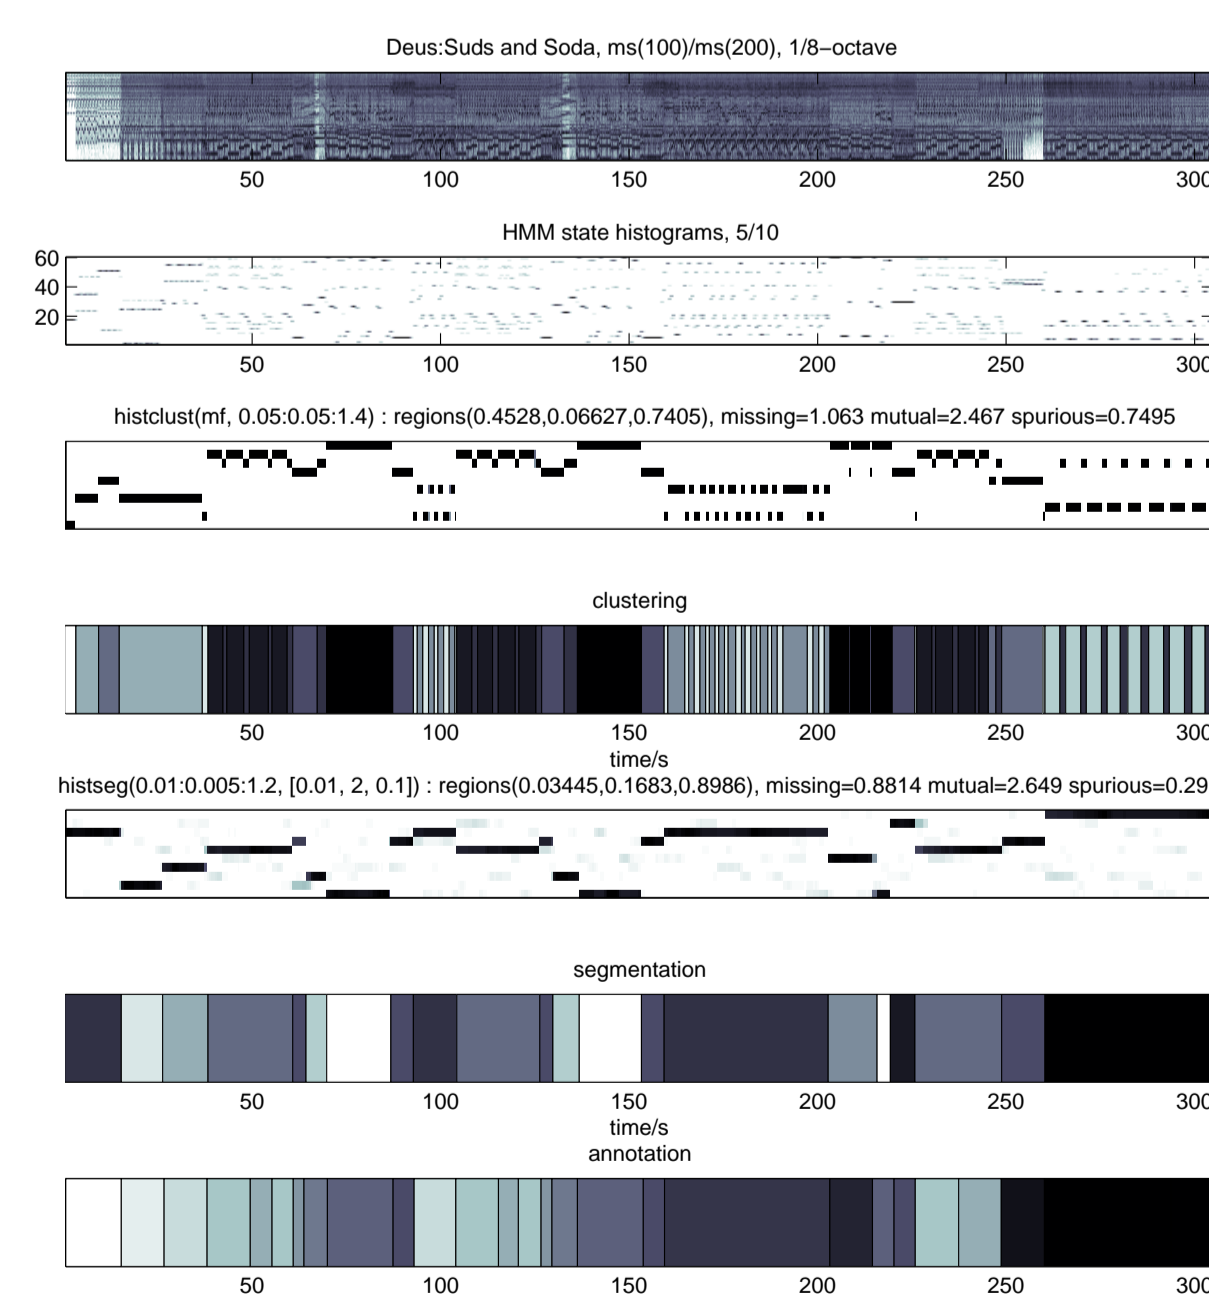
The Spice Girls: Wannabe



A Ha: Take On Me



Deus: Suds and Soda



The Clash: Should I Stay or Should I Go

