

Temporal and Top-Down Processing in Audition



Malcolm Slaney
Yahoo! Research
Stanford CCRMA



Outline


- The Non-Linear Cochlea
- Correlogram
- Pitch
- Information Flow
- CASA
- Saliency
- Relational Networks

Cochlear Nonlinearity




Hair Cell Motility

Cilia




Pipette (changing membrane potential)


Auditory Nerve Patterns



CHARACTERISTIC FREQUENCY, KHZ




Correlogram Architecture




Tone Correlograms


200Hz Tone



Harmonics of 200Hz




Tone Examples




With help from Richard O. Duda

Speech Correlogram

Unvoiced /s/




Voiced Vowel



Speech Examples

- Speech Examples
 - Leonardo at Apple





RTCorrelogram


21705.12

8899.67

3372.62

1143.70

35.33



Input Type: Default




Display Type: Correlogram




Cochlear Downsample Factor: 2 (22050.00)

Print Current Stats


Pitch





Pitch Architecture



Pitch in Noise




Ambiguity of Pitch




Pitch Theories


A "Correlogram"
*Temporal implementations with
neural delay lines*



B *Central Pattern Match Theory*



Correlogram Implementation



Today

FFT

Licklider

Neural Delay

Patterson

Strobed-
temporal
integration


Alternative

Oscillators

AM Modulator


Processing Chain

**One Dimensional
(waveform)**




Cochlear Processing

**Two Dimensional
(not a spectrogram)**





Correlogram Processing

**Three Dimensional
(neural movie)**



Block Diagrams



Problems



Problems

- Continuation
 - Tone and Noise
 - Parliament Cough



Problems

- Continuation
 - Tone and Noise
 - Parliament Cough
- What do you hear?
 - Waveforms?
 - Ideas?



Speech Examples

Wedding



Sine



Natural



Visual Context

TAE

CAT

What Vowel is This?



Peter Ladefoged




Word 1





Word 2



Word 3












McGurk




















Information Flow




ASR


Language model for
the words: “one”,
“two”, “three”





Word model showing
phonemes for the word
one



Acoustic (phoneme) model
for the phoneme /ʌ/



Streaming of Tones



Auditory Scene Analysis



All harmonics - Oboe


Even harmonics - Soprano

Odd harmonics - Clarinet

Old plus New Principle

The effect of context

- Context can create an ‘expectation’:
i.e. a bias towards a particular interpretation
- e.g. Bregman’s “old-plus-new” principle:
A change in a signal will be interpreted as an *added* source whenever possible




- a different division of the same energy depending on what preceded it

Conventional Scene Analysis

Computer models of grouping



- “Bregman at face value” (e.g. Brown 1992):




- feature maps
- periodicity cue
- common-onset boost
- resynthesis

Goto—CASA with MIDI


MIDI Sequence



Barker—ASR




Ellis—Prediction Driven



Saliency

A




Mechanisms for Allocating Auditory Attention: An Auditory Saliency Map

Christoph Kayser,* Christopher I. Petkov,
Michael Lippert, and Nikos K. Logothetis
Max Planck Institute for Biological Cybernetics
Spemannstrasse 38
72076 Tübingen
Germany

conceptualized as theoretical models, w
hierarchical and parallel extraction of d
and build on existing understanding o
cessing. For the visual system, such
shown to replicate several properties
attention [2, 8, 9]. For other sensory sys


Saliency Example

- Time-frequency display
- Saliency map shows high-interest locations



Saliency Maps

- Longer tones better
- Missing parts salient
- Modulation more salient
- Forward masking works



Sound Examples

- Birds



- Calls



- Cows



- Horse




- Waterfall




Saliency Comparison

- Details of saliency comparison
- Model predictions




Relational Network (Simple)

- Patches of neurons
- Each measure one quantity
- Bidirectional relations for feedback/feedforward




Thanks to Rodney Douglas

Relational Network (example)



ASR Relational Network




Desired Results


Relational Feedback
With




Desired Results



Desired Results




Desired Results




Simulation

Iteration 1




Simulation 2

Iteration 2




Simulation 3

Iteration 15




Grossberg— ART



The LAMINART model. From Grossberg, S. (1999). How does the cerebral cortex work? Learning, attention and grouping by the laminar circuits of visual cortex. *Spatial Vision*, **12**, 163-186

Statistical Means

- ICA
 - Different distributions
- One Microphone
 - GMM models of distribution



Separated signals after 4 steps of FastICA

Semantic-Audio Retrieval



Semantic-Audio Retrieval



- Predicted words:
 - anim (.108), hors (.105), left (.086), trot (.065), approach (.059), track (.047), walk (.040), depart(.037)



Semantic-Audio Retrieval



- Predicted words:
 - anim (.108), hors (.105), left (.086), trot (.065), approach (.059), track (.047), walk (.040), depart(.037)
- True label:
 - animals: horses, two horses trot past on rough track left to right



Semantic-Audio Retrieval



- Predicted words:
 - anim (.108), hors (.105), left (.086), trot (.065), approach (.059), track (.047), walk (.040), depart(.037)
- True label:
 - animals: horses, two horses trot past on rough track left to right



- Predicted words:
 - bird (.11), ambienc (.107), jungl (.104), morn (.094), Africa (.093), anim (.054), bark (.029), dog (.020), cricket (.018)

Semantic-Audio Retrieval




- Predicted words:
 - anim (.108), hors (.105), left (.086), trot (.065), approach (.059), track (.047), walk (.040), depart(.037)
- True label:
 - animals: horses, two horses trot past on rough track left to right



- Predicted words:
 - bird (.11), ambienc (.107), jungl (.104), morn (.094), Africa (.093), anim (.054), bark (.029), dog (.020), cricket (.018)
- True label:
 - jungle, Africa, Africa: morning ambience, birds


Audio Anchors

- Process
 - Build models for known sounds (GMM/HMM)
 - Measure likelihood of new sound
 - Form vector
- Advantages
 - Summarizes entire sound
 - Agglomerative clustering using distances




Audio Features

- MFCC/Rasta at 100Hz
 - 13 cepstral coefficients per frame
- Stack frames across time
 - +/- 3 frames
- Linear discriminant analysis (LDA)
 - 10 classes guide best feature decision



Model Clustering Algorithm

- E-Step
 - Build model for all signals in cluster
 - GMM (ignores time dependence)
 - HMM (considers time dependence)
- M-Step (for each signal)
 - Measure likelihood for all models
 - Assign to model with highest likelihood



Multinomial Distribution

- Bag of words model
- Likelihood of document given model

$$L = \prod w_i^{n_i}$$

- w_i is probability of word i
- n_i is count of word i
- L is likelihood of document



Semantic Clustering

- Find related documents
- Multinomial EM (Nigam)
 - Probabilistic model of word probabilities
- Include smoothing in model

Document 10

anim (0.2)
cattl (0.2)
farm (0.2)
restless (0.2)
shed (0.2)
abrupt (0)
abruptli (0)
abuko (0)
acanthi (0)
acoust (0)

Cluster 383

anim (0.125)
sheep (0.125)
exterior (0.0833333)
past (0.0833333)
run (0.0833333)
tree (0.0833333)
wind (0.0833333)
concret (0.0416667)
flock (0.0416667)
gate (0.0416667)


Cluster 388

anim (0.137931)
gallop (0.137931)
hors (0.137931)
pass (0.103448)
singl (0.103448)
dirt (0.0689655)
gravel (0.0689655)
concret (0.0344828)
folei (0.0344828)
group (0.0344828)


Words and Image Spaces

Word Space

Image Space




Words and Image Spaces



Words and Image Spaces


Word Space

Image Space



Mixture of Experts

- $\mathbf{x}^{(n)}$ is input query
- Different experts are good with different kinds of queries
- Gate chooses expert
- $y^{(n)}$ is classification
- Train with EM
- Want probability estimates



MPE Mathematics


$$P(a | q) = \sum_c P(a | c) P(c | q)$$

Diagram illustrating the MPE equation with semantic annotations:

- $P(a | q)$ is annotated as **Prob(audio|query)** (indicated by an arrow).
- $P(a | c)$ is annotated as **Prob(audio|cluster)** (indicated by an arrow).
- $P(c | q)$ is annotated as **Prob(cluster|query)** (indicated by an arrow).


(Mixture model, interpolation)

MPE Mathematics




Semantic Retrieval

- Acoustic to Semantic
 - Train multinomial per cluster
 - Label a sound




Acoustic Retrieval

- Semantic to Acoustic
 - Train semantic multinomial per cluster
 - Label a sound




Training Overview

Audio Processing





Text Processing




Test Results


- Semantic → Acoustic
 - Median rank is 9
- Acoustic → Semantic
 - Median rank is 17.5.








Conventional




Conventional




Better?



Better?



Hadoop



Thanks

malcolm@ieee.org