Nonlinear Spectro-Temporal Integration in Ferret Primary Auditory Cortex

Ivar Thorson

Laboratory of Brain, Hearing, and Behavior
David Lab

May 16, 2014

Confession: I'm not a neuroscientist

2004	B.S. Electrical Engineering *University of Washington*
2008	M.S. Mechatronics *Nagoya University, Japan*
2012	Ph.D. Advanced Robotics *Istituto Italiano di Tecnologia, Italy*
2012-	Neuro-analyst for Prof. Stephen David *Oregon Health & Science Univ.*

WARNING:

Stephen David *et al.* are not responsible for blatant ignorance of neuroscience I may display during this presentation.

DANGER:

I also promised my wife Daniela that I wouldn't embarass her in front of all her colleagues.

Engineering Method: Input-Output Models

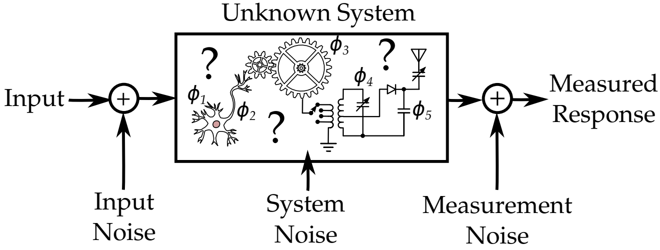

- Given input and output data, we want to identify:
  - the *most important structures* inside the "black box"
  - the values of the *model parameters* $\phi_1, \phi_2, \ldots$
  - the time-varying *state* (voltage, velocity, etc)
  - as accurately as we can despite *uncertainty/noise*
- Engineering methods are applicable to *any input-output system*:
  - **Mechanical**: Swinging doors, engines, solar systems
  - **Electrical**: Single resistors, circuits, computers
  - **Biological**: Neurons, stimulus→response sensory systems

Input-Output Models of the Auditory Brain

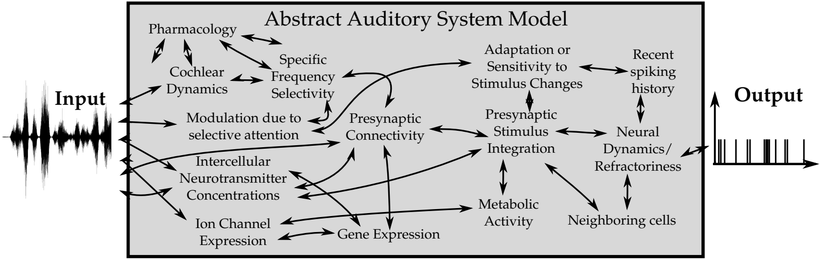

Goal: Use engineering methods to develop a complete functional model of the black box spanning from the ear to the auditory cortex.

- We want models that map sounds to predictions about spiking
  - Modeling every detail is too difficult
  - Which physiology is *most important/relevant*?
  - Better predictive models $\leftrightarrow$ better understanding
- Understand brain computation at an *algorithmic* level[1]
- *Technological applications* in cochlear implants, hearing aids

[1] Marr 1982

How are models useful to physiologists?

1. *Models give researchers flexibility in their stimuli.*
   - Tuning curves easily estimated with simple pure tones...
   - ...but predicted responses extrapolated from simple stimuli can poorly match responses to complex/natural sounds [2]
   - We infer tuning from neural responses to *natural stimuli*
2. *Models can be used post-hoc to "data-mine" experimental data*[3]
   - Test multiple new hypotheses on old data
   - Select the best model
   - Model parameter values are contextual measurements
3. *Models can hint at future experiments.*
   - If we notice clusters of model parameter values, can we categorize neural types from functional properties?[4]

[2] Theunisson 2000, David 2009

[3] Mesgarani 2014

[4] Woolley, 2009

Approach Overview

- We use natural stimuli and awake animals
- We use physiology to motivate mathematical terms
- We test many, many alternative models on the same data set(s)
  - Published and unpublished models
  - Many combinations of model terms[5]
- We quantify how much each model term helps

[5] In the last 18 months, we have fit >540,000 models

What system is the David Lab studying?

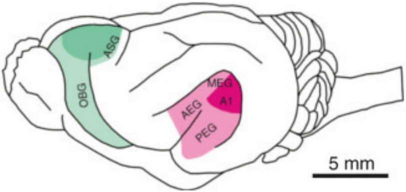

- Today's data is from the ferret
  - Hearing range overlaps humans (20Hz-40kHz)[6]
  - Network of well-defined auditory cortical areas[7]
  - Can be trained for behavioral experiments [8]
- ...but data from any animal can be used (mouse, marmoset)
- Today's data is from primary auditory cortex (A1)
- ...but data from other sensory regions can be used

[6] Kelly, 1986
[7] Bizley 2005
[8] Fritz 2003; David 2012; Bizley 2013

How are we acquiring/analyzing data?

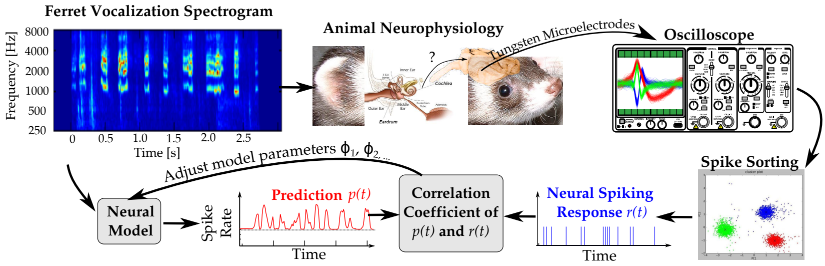

We record single-unit activity:

1. Play a sound (Above: ferret vocalizations in spectrogram form)
2. Record extracellularly with tungsten micro-electrodes
3. Spikes are isolated, sorted as single units via PCA
4. Estimate model parameters from stimuli and spikes
5. Evaluate model performance using novel data with correlation coefficient.

Spectro-Temporal Receptive Field (STRF)

Imagine playing many different infinitely brief pure tones...

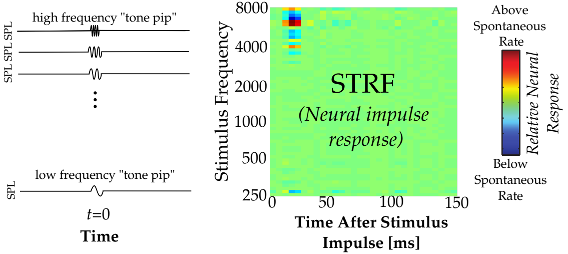

...and recording when/how the neuron responds to each "pip"

The STRF[9] can predict auditory neural responses to any stimulus:

1. Break a stimulus into many brief "tone pips"
2. Find the response to each tone pip with the STRF
3. The prediction is the sum of the responses to all pips

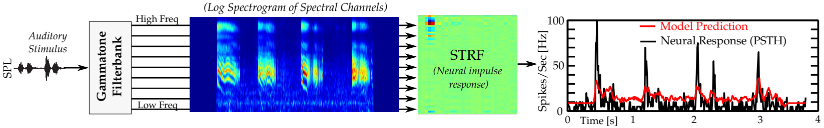

[9] Aertsen 1982, deCharmes 1998, Theunissen 2001

Models as Measurements

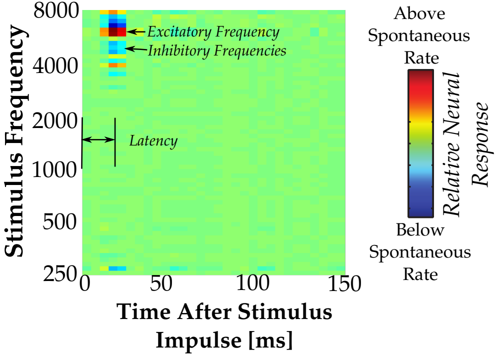

The STRF model as a "measurement" tells us:

- The frequencies to which the neuron is sensitive
- The latency between stimulus and neural response
- "Stationary temporal dynamics": onset-sensitive or integrating
- Possible sensitivity to harmonicity, frequency sweeps, etc.

What's wrong with the STRF?

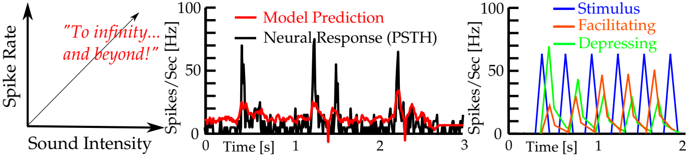

The STRF makes obviously erronious "linear" predictions:

1. Doubling sound intensity *always* doubles the spike rate, and neurons have *no limit* to how fast they can fire
2. *Negative* spike rates are possible
3. Neural responses cannot exhibit *depression* or *facilitation*
4. Non-additive interaction between frequency bands is not modeled

We are interested in correcting these deficits.

Improvement 1/4: Volume Compression

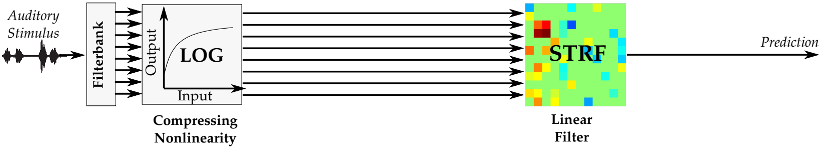

Observations: Neuron spike rates often respond logarithmically to increased sound intensity

Improvement: A base-n logarithmic compression term [10]

$$f_1(s) = \log\left(s + \phi_1\right)$$

where s is the input.

- One parameter: ϕ_1
- Improves predictions by 11-15%

[10] Gill & Theunisson 2006

Improvement 2/4: Base and Max Thresholds

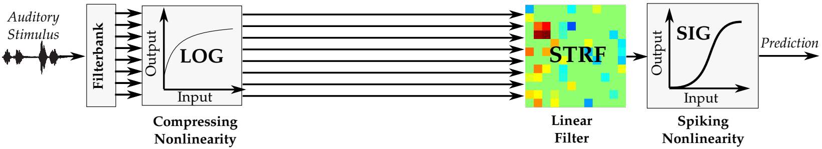

Observations:

1. Very weak stimuli may not stimulate neurons
2. Very strong stimuli may saturate the neuron at a max firing rate

Improvement: Asymmetric logistic sigmoid spiking term[11]

$$f_2(s) = \phi_2 + \frac{\phi_3}{1 + e^{-\phi_4(s-\phi_5))}}$$

- Five parameters: base rate, max rate, center inflection point, low side curvature & high side curvature
- Improves predictions by additional 7%

[11] Nykamp & Ringach 2002.

IMPROVEMENT 3/4: NONLINEAR DYNAMICS

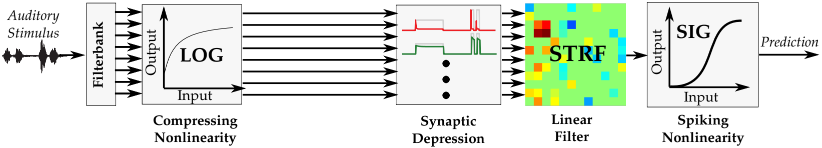

Observations: Neurons may respond more weakly for a few moments following strong stimuli

Improvement: Model time-varying "synaptic depression" [12]

$$\dot{d}(t) = -\phi_6 \cdot s(t) \cdot d(t) + \frac{1 - d(t)}{\phi_7}$$

$$f_{DEP}(s) = s(t) \cdot d(t)$$

- A name is just a name – could be local feedback inhibition
- Two parameters: depression rate and recovery rate
- Improves predictions by additional 6-7%

[12] Markram 1998, David 2013

Improvement 4/4: Multiplicative Interaction

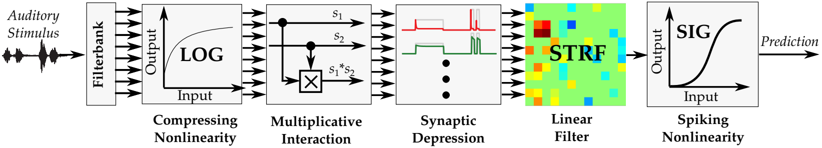

Observations:

1. Neurons may respond weakly to stimulus s_1 alone
2. Neurons may respond weakly to stimulus s_2 alone
3. Neurons may respond **strongly** to both stimuli simultaneously

Improvement: Model multiplicative terms.[13]

- Linear Only:

$$p(t) = \phi_1 s_1(t) + \phi_2 s_2(t)$$

- Linear + Multiplicative:

$$p(t) = \phi_1 s_1(t) + \phi_2 s_2(t) + \phi_3 s_1(t) s_2(t)$$

- Improves predictions by additional 4-5%

[13] Eggermont, 1993

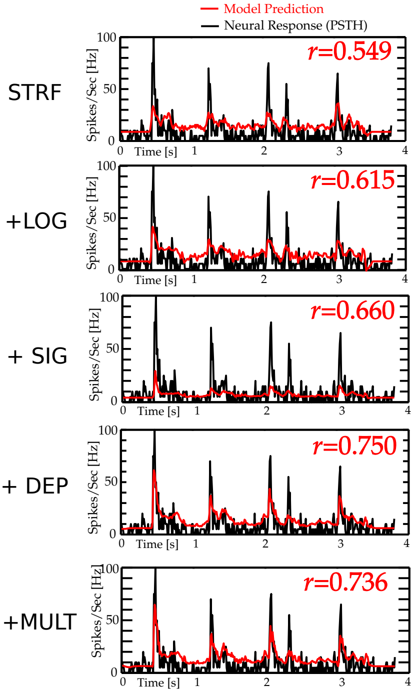

Model Prediction
Neural Response (PSTH)
STRF
r=0.549
+LOG
r=0.615
+ SIG
r=0.660
+ DEP
r=0.750
+MULT
r=0.736
Spikes/Sec [Hz]
Time [s]

Mean Prediction Correlation

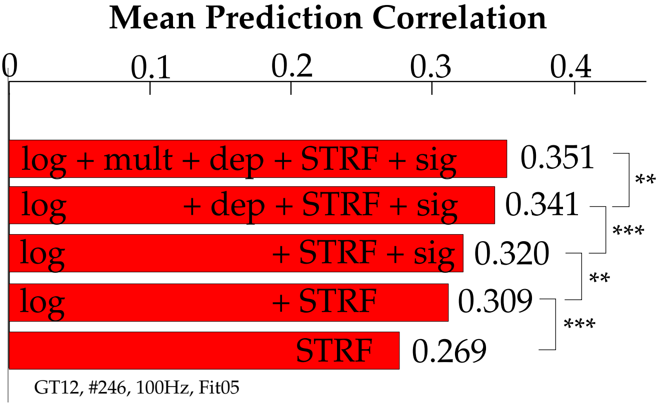

- Tested across a population of N=167 neurons
- Improvement of mean performance: 31%
- Best single unit prediction correlation to date: 0.9082

Big Picture: Performance

- Some neurons we can describe very well
- Others we cannot predict
  - Some not strongly driven by sound stimuli...
    ...(maybe we'll find models that explain them later?)
  - Using same model for *every* neuron in cortex
- The more synapses from the cochlea, the harder to model

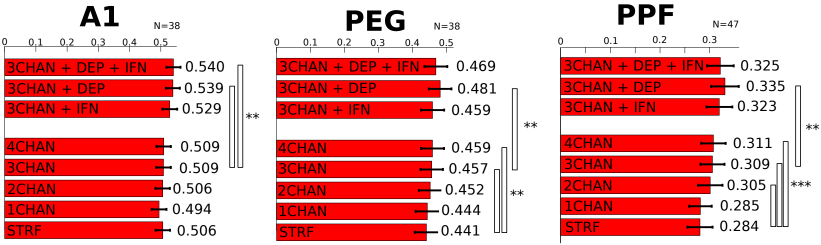

r_ceiling values, r_val < r_floor cut out, gt12, Batches #264, 265, 266, 100Hz, Compressor + NL, fit05

- Good models for A1 seem applicable later
- Would earlier, more "peripheral" areas be easier?

Big Picture: Complexity

"Are our more complex models doing better just because they have more parameters and could describe any data set better?"

- We use fresh, novel data for measuring model performance, so overfitting cannot occur
- We work hard to reduce the number of parameters using matrix factorizations and re-parameterizations
  - Fewer parameters are more comprehensible
  - Models with fewer parameters require less data
  - Particularly useful in behavioral studies

Parameters / Performance Tradeoff

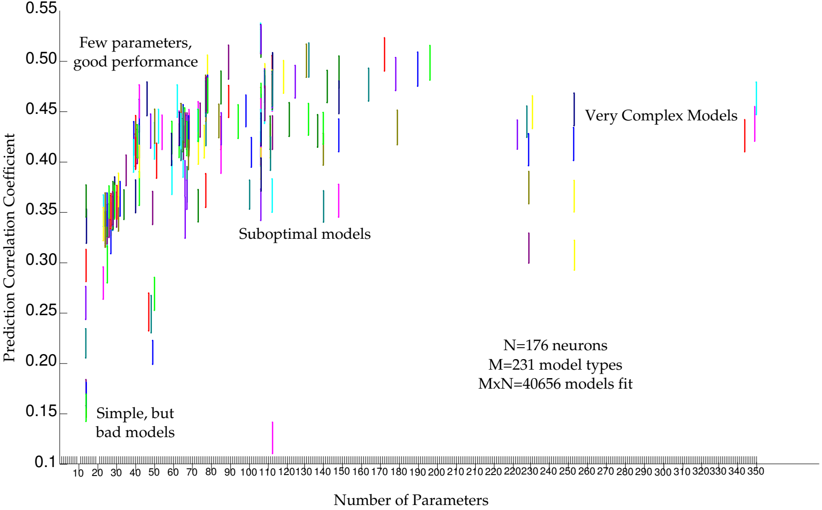

Each bar shows a different model's mean +/- standard error of prediction correlation across a population of 176 neurons.

Conclusions

- Predictive models can help *you* learn more from *your* experimental data (even post-hoc)
- Classic models like the STRF have some shortcomings
- Incorporating biologically-inspired functional terms helps
- Most hypotheses out in the literature are sub-optimal
- We improved performance of the state of the art by 30%
- We have greatly reduced the number of parameters needed

Thanks to

- Stephen David
- Sean Slee
- Jean Liénard
- Daniela Thorson
- Henry Cooney
- Brian Jones
- Zachary Schwartz

Questions?

STRF Revisited

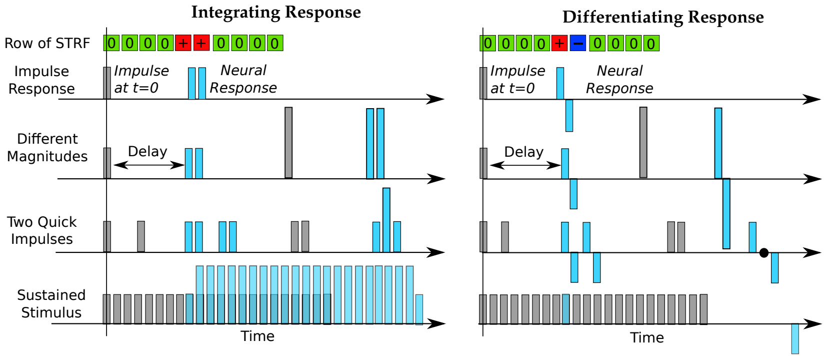

When is Model A is better than B?

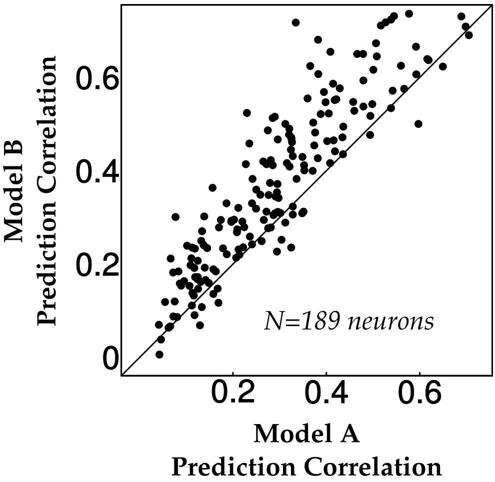

- There is considerable natural variation in neural function
  - Model A wins for some cells
  - Model B wins for other cells
- Does that mean A and B describe different cell types?
- Comparing parameters from different models is hard
- We'd prefer the best possible "common yardstick"

Questions Considered

1. Which combinations of nonlinearites are best?
2. Does incorporating better models of the cochlea improve prediction scores in A1?
3. Do models fit using alternative performance metrics differ qualitatively?
4. Do assumptions of smoothness or sparsity help?
5. Which optimization algorithms are the best for neural data?
6. How close are we to the upper perforamnce bound possible?

Math in Matrix Form

Showing only 1 channel of 12-36 channels, showing only the 1st-order filterbank instead of a 4th-order filterbank, and ignoring multiplicative interaction:

$$
\begin{aligned}
\begin{bmatrix} \dot{x}_{p1z1} \\ d_{out} \\ \dot{d} \\ y_{logn} \\ \dot{x}_1 \\ \dot{x}_2 \end{bmatrix}
&= \begin{bmatrix} -p_{IIR} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{-1}{\phi_5} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & R & S \\ 0 & 0 & 0 & 0 & -S & R \end{bmatrix} \mathbf{x} \\
&\quad + \log\left(\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & K\frac{R-wz}{S} & K \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \mathbf{x} + \phi_1 \right) \\
&\quad + \begin{bmatrix} 0 \\ ds \\ \frac{1}{\phi_5} - \phi_6 ds \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} u(t-\phi_4) \\
y_{IIR} &= \begin{bmatrix} K_2(z_{IIR} - p_{IIR}) & K_2 & 0 & 0 & 0 \end{bmatrix} \mathbf{x} \\
\mathbf{y} &= \phi_0 + \phi_1 e^{-e^{\phi_2 (y_{IIR} - \phi_3)}}
\end{aligned}
$$

Extra Mathematical Comments

- Neuronal activity is spectrally complex but temporally simple
- Weighting and summing is a great method for approximating any function and is neurobiologically plausible
- The best nonlinearities use the natural number e
- Regularization and smoothing almost never helps
- Our model can be efficiently computed as a system of nonlinear differential equations
- Developing better spectral basis functions is very hard
- The performance metric and fitter are equally as important as model structure.