Looking for taste

Shane Hoversten

Why do we like what we like?

Biederman's findings with scene evaluations

- Some innate preferences: safety, good sight lines, etc.
- But also, opioids in para-hippocampal region
- So it's pleasurable to see stuff you can make sense of
- This is an individual metric!

More specifically:

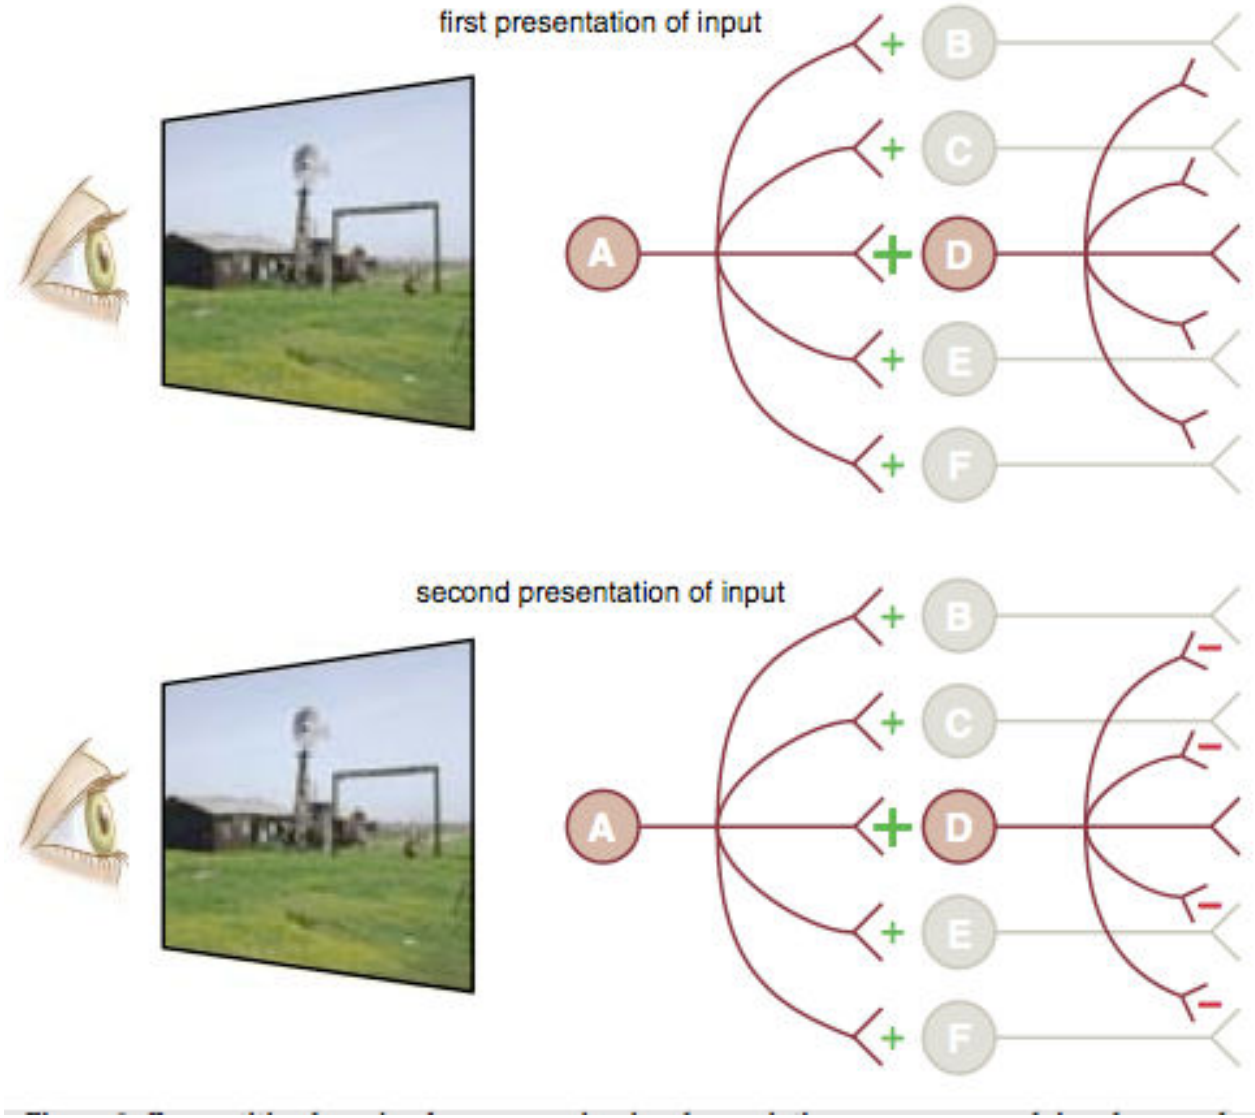

Something like this…

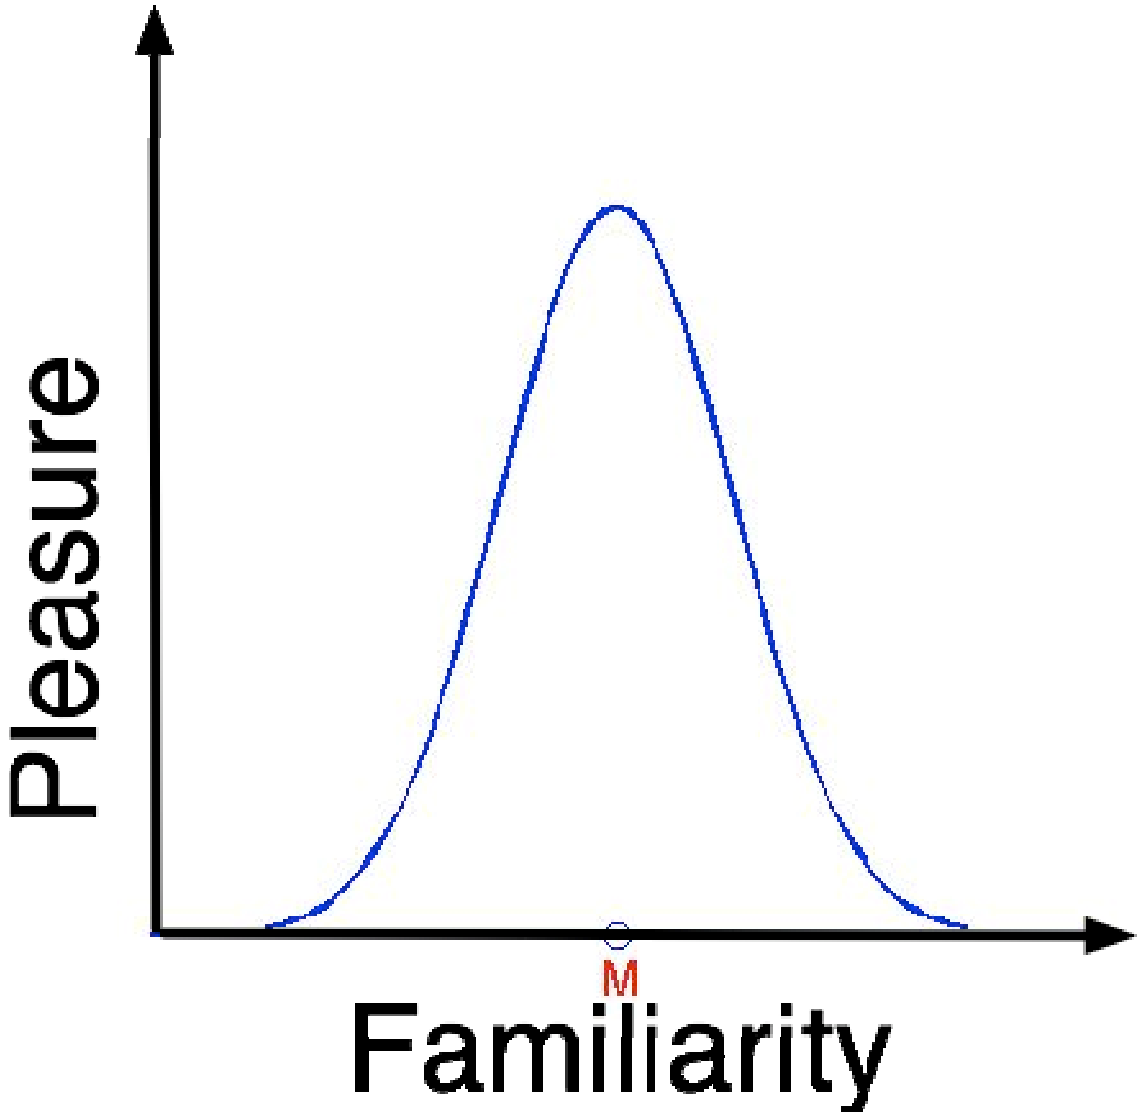

Look familiar?

The “oh!” of pleasure

- The more neural activity in this region, the more opioid activity, the more pleasure
- However, stuff has to “make sense” in order to activate neural connections through PH
- Again, the same “stuff” can mean different things, and give different amounts of pleasure
- I call this “informatia.”
- For example:

What's this?

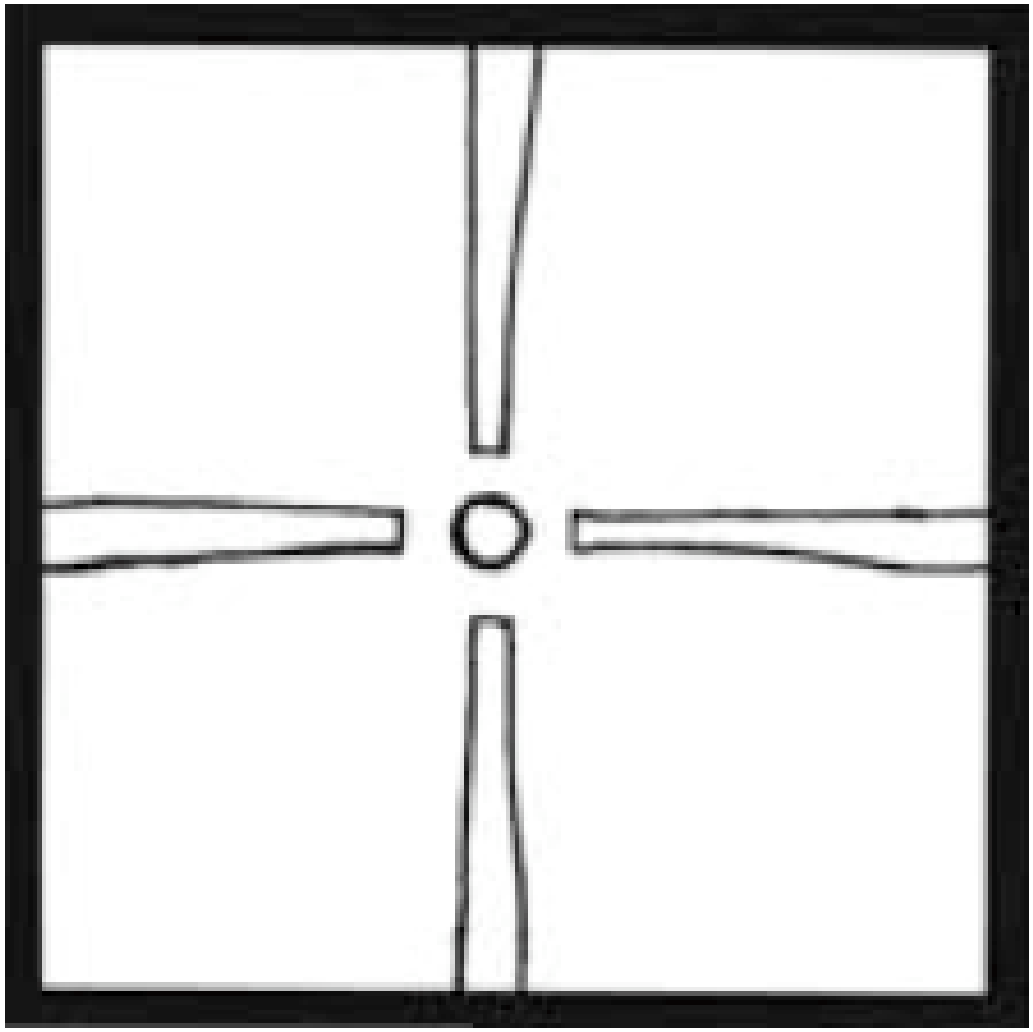

Four elephants smelling an orange

How about this?

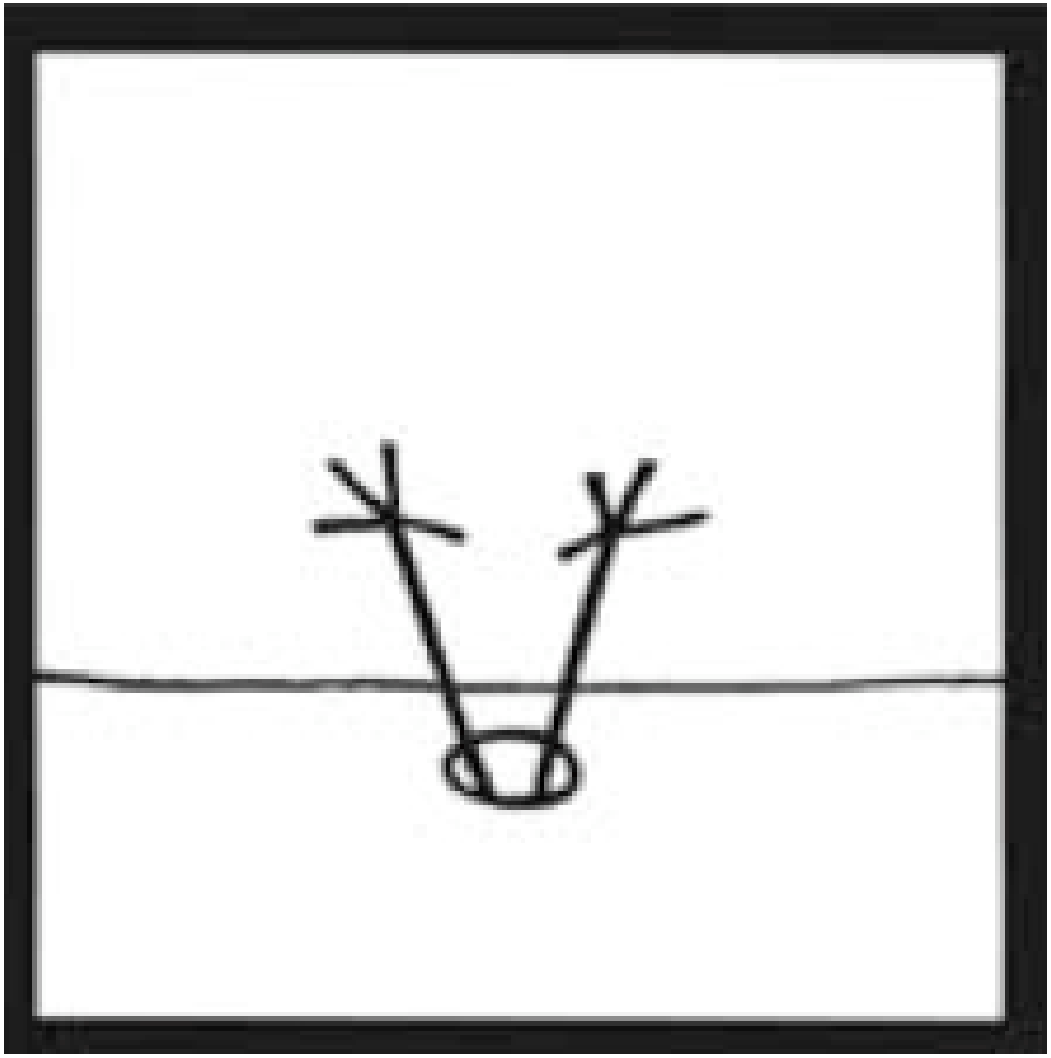

An early bird encountering a very strong worm

What's going on?

- Adding a new fact (the answer) increases the informatia, increases the opioid response.
- What are the implications for music?
- Check out how fantastic this is:

The idea

In a nutshell, this is the hypothesis:

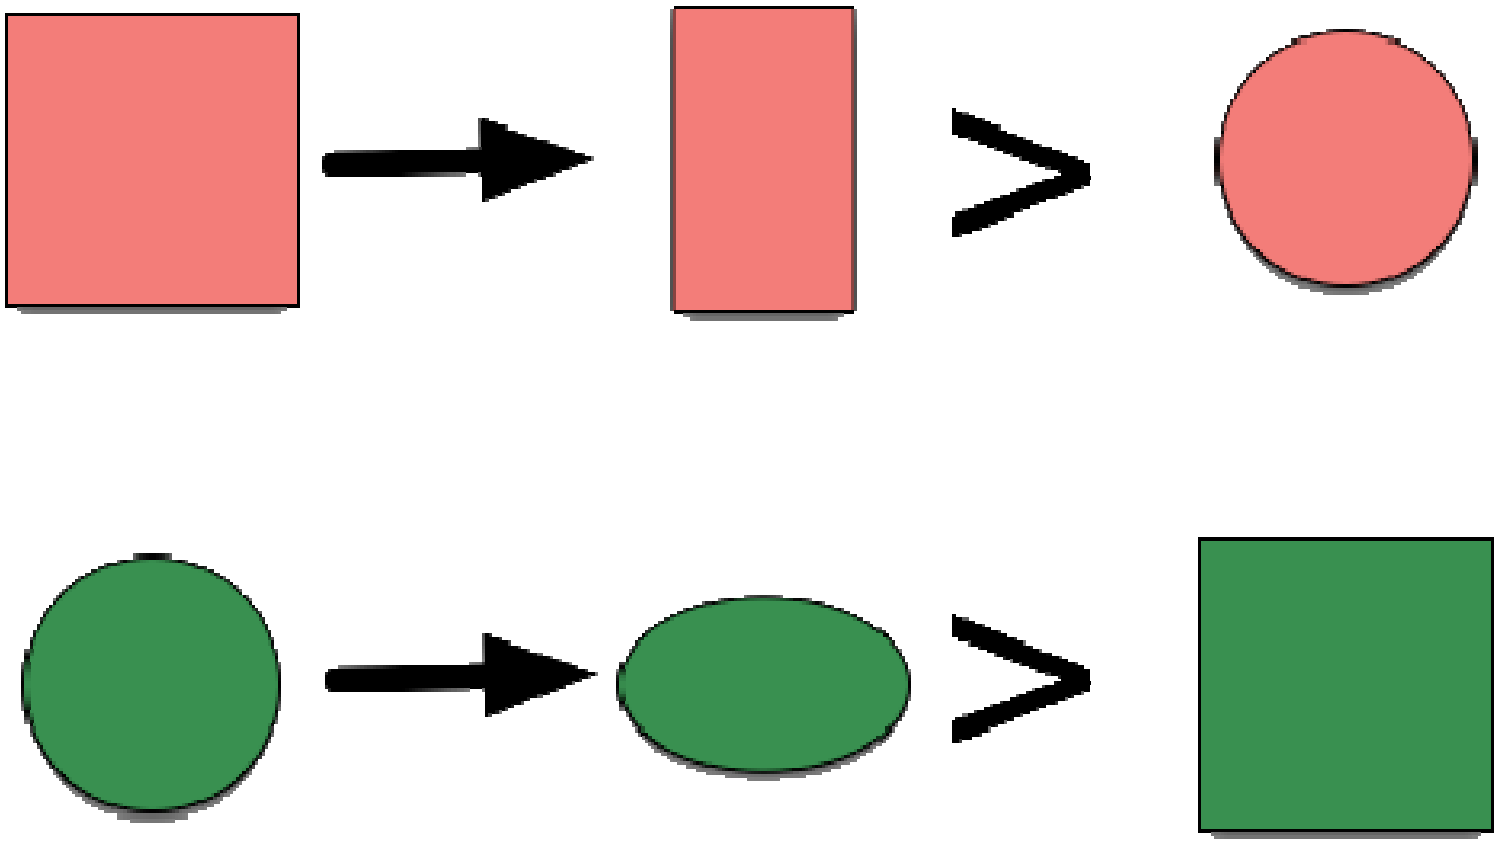

Methodology

- Take a bunch of midi files of type x
- Build a generator for stuff like x: g(x)
- Do the same thing for y
- Now tweak g(x) to be g'(x)
- Do the same thing w/ y
- Subjects trained on g(x) should prefer g'(x) to g(y)
- In fact, we should be able to predict HOW MUCH they prefer it. ("should"…)

Real methodology

- Didn’t do any experimental stuff
- Instead, try to make a framework that could make possible these experiments
- There’s a lot of complication here that I won’t talk about (e.g., the issue of pre-existing taste skewing the experiment.)

Generator

- Use factor oracle
- After much screwing around, wound up re-creating Alex's framework to deal w/ difficulties in representing time
- Collapse all midi channels onto single channel (which does weird things to multi-channel music)
- Tune-able replication/recombination
- Python/Lisp hybrid

Methodology

- Search for “ethnic” midi files
- Extract midi, combine together
- Feed to oracle to extract info, spit out generator for “ethnic” music
- Here’s Japanese: 
- Here’s Jewish folk: 
- Here’s Japanese w/ a taste of folk: 

Issues

- Multi-channel midis make generated music unnecessarily “busy”
- “Event equivalence” should probably be smoothed out - binned, maybe.
- For these reasons, seems most coherent with fewest inputs.

How to measure information differences in music?

The idea of “runs”

- A run is a sequence of notes
- Can be recognized by a generator or not
- Consequently, there are positive and negative runs
- Runs vs. generators seems like a candidate information-gain metric…

Japanese-Japanese

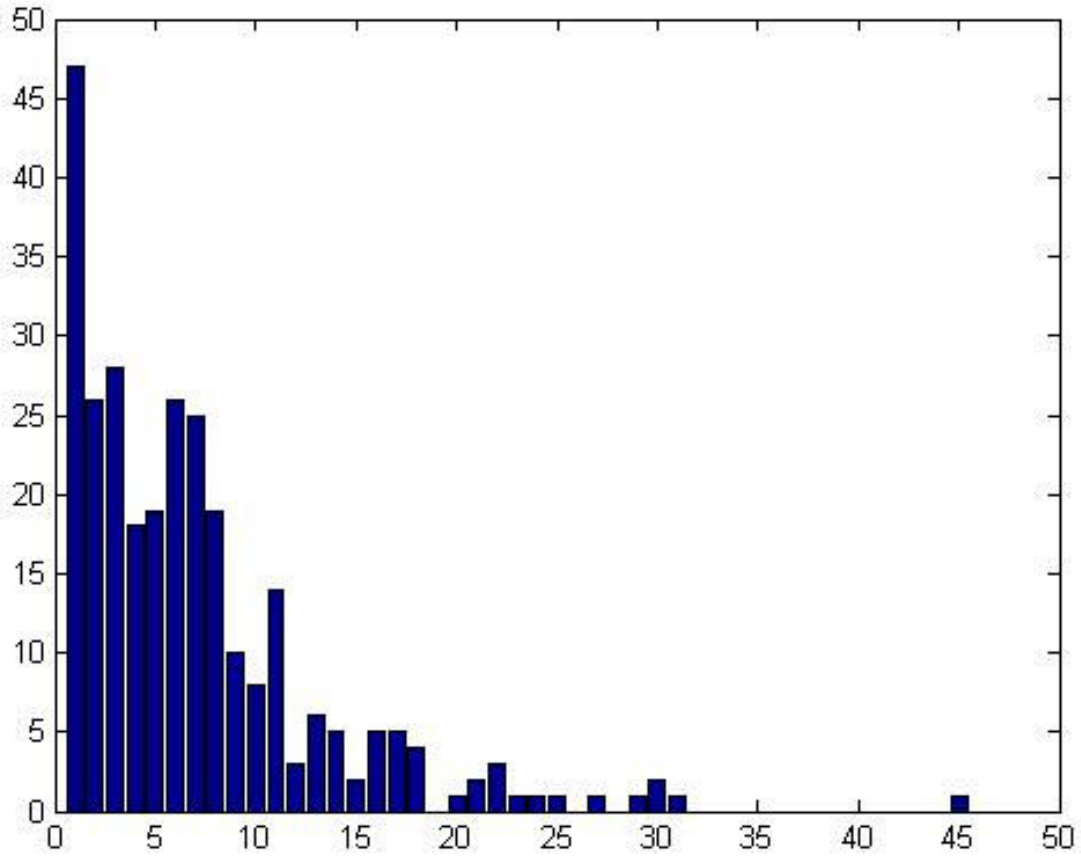

Jap.-Fus and Jap.-Jewish

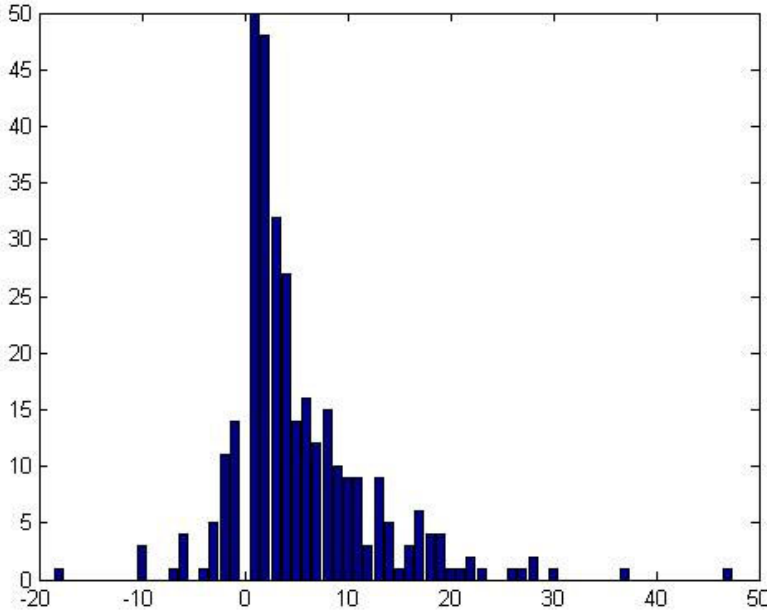

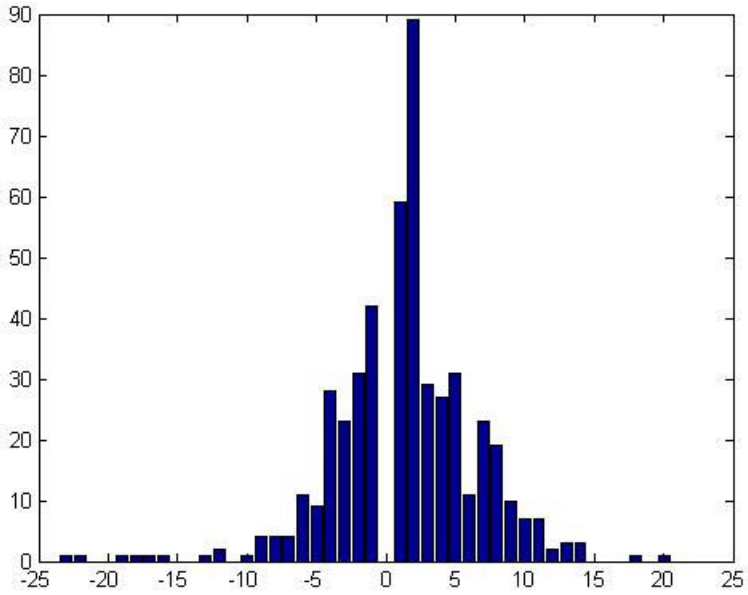

Jewish-Jewish

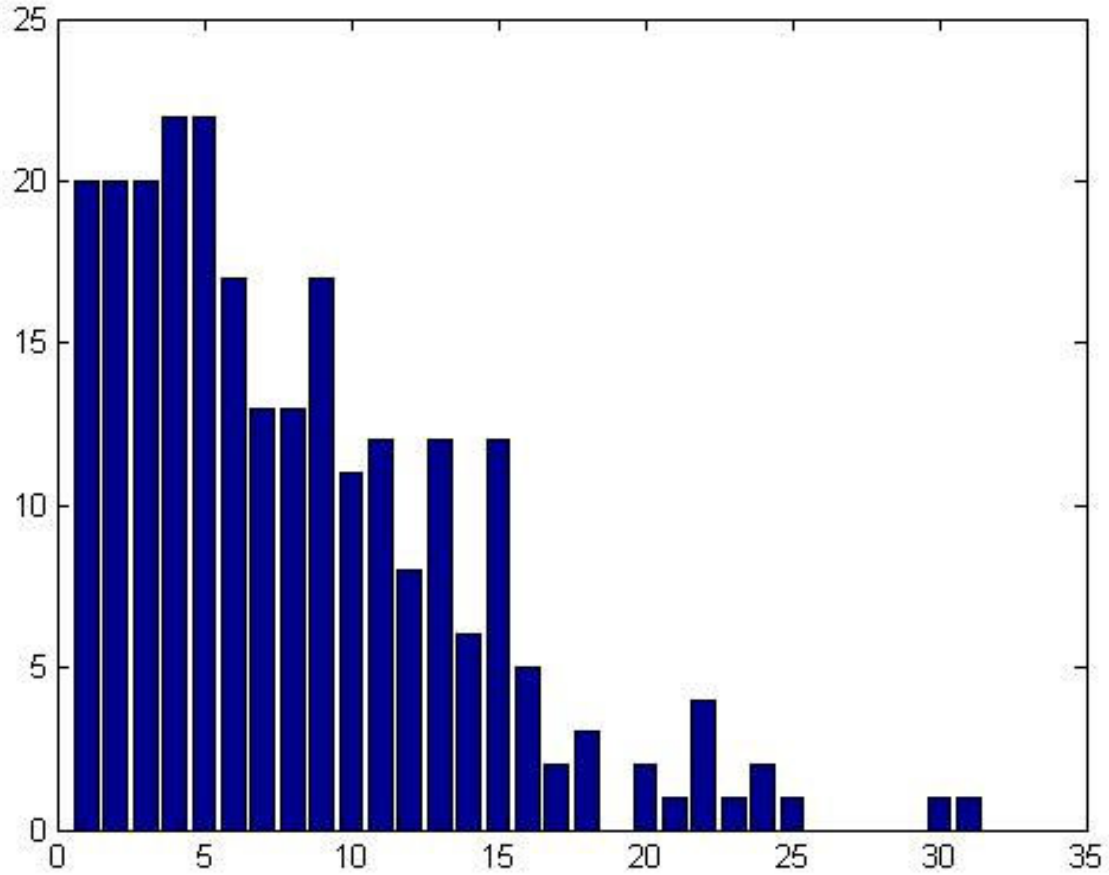

Jewish-Fus., Jewish-Jap.

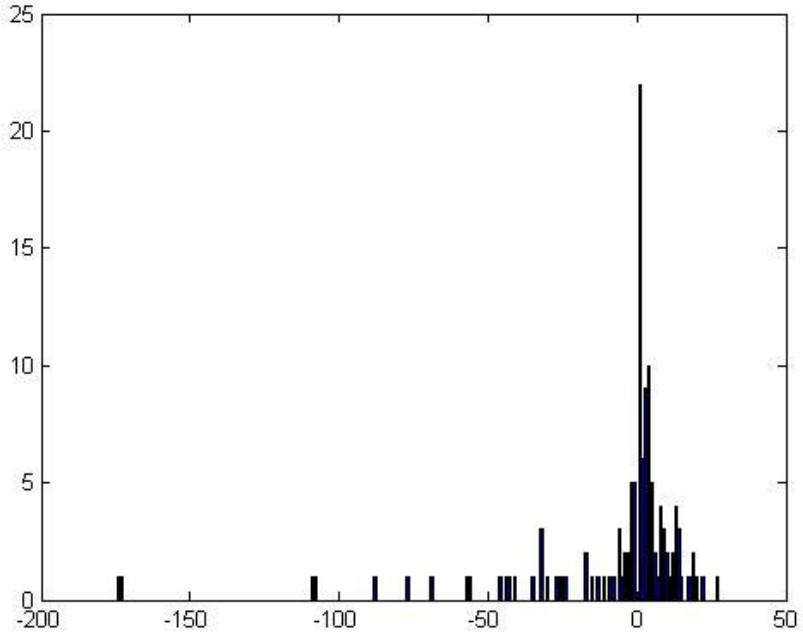

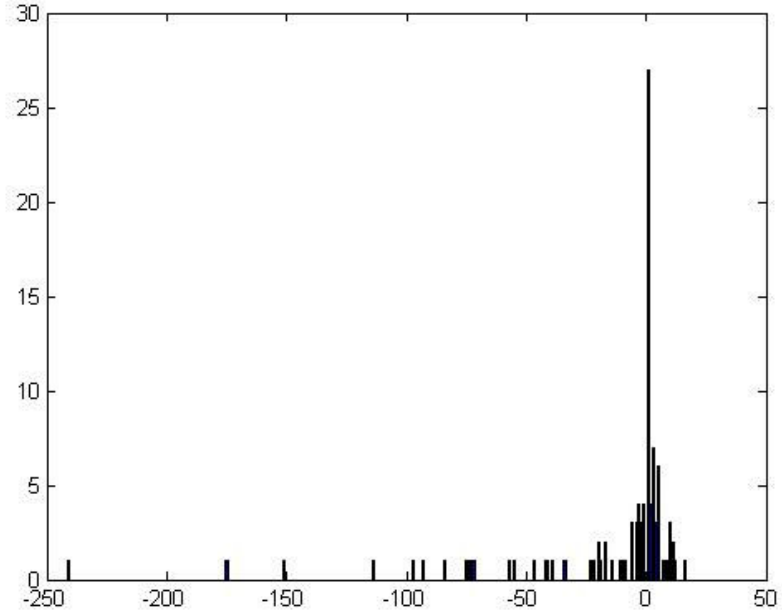

Fusion-Fusion

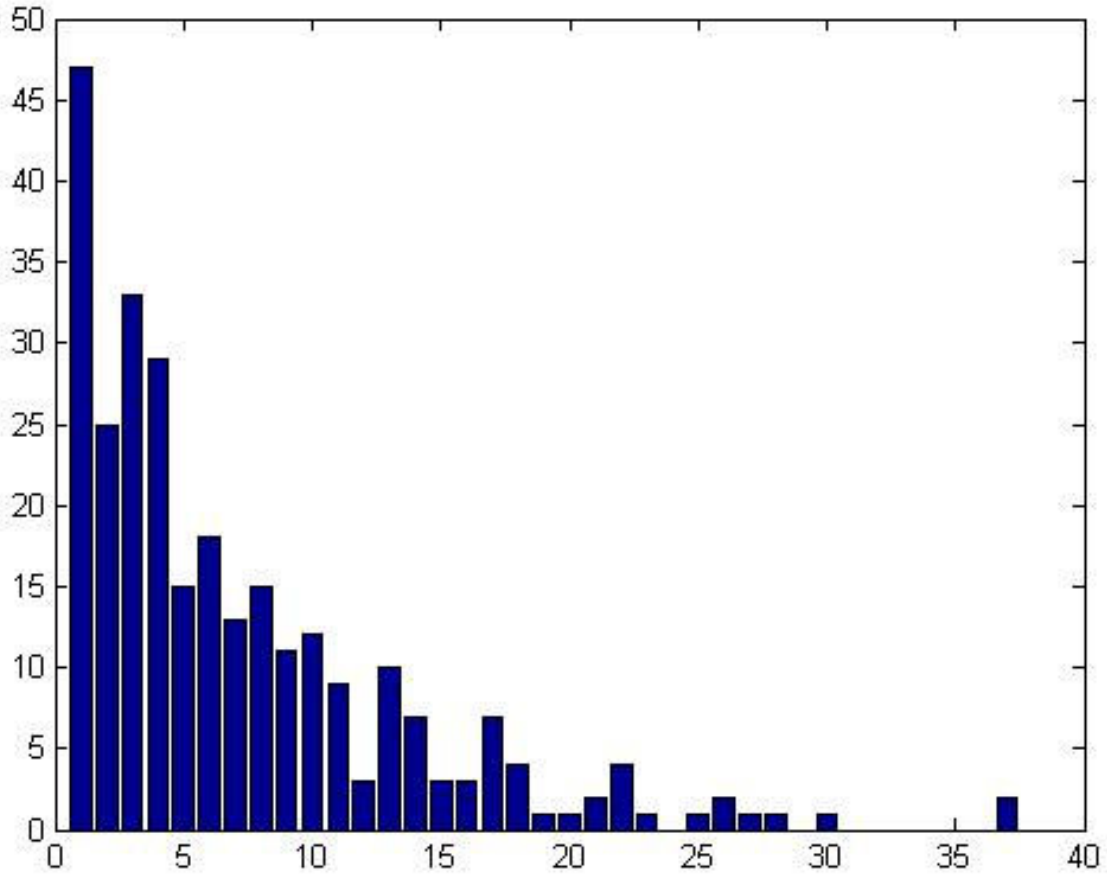

Fus-Jap., Fus-Jewish

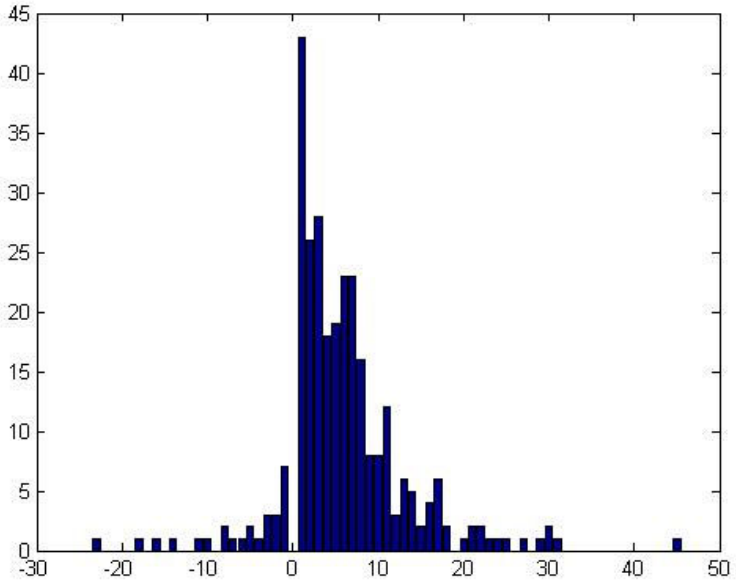

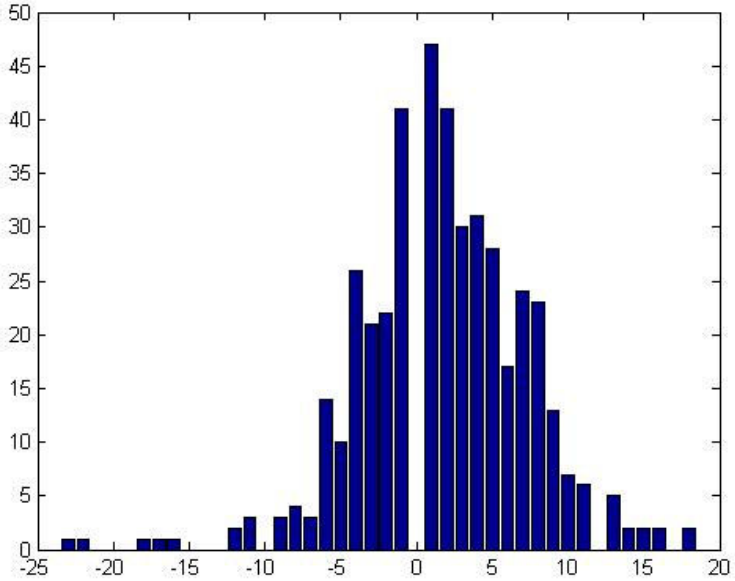

Uhhh …

That left something to be desired

How to measure informatia?

- To do the most interesting part of this experiment, we need a better way to measure informatia.
- I don’t know how to formulate this idea, even in the simplified domain of simple tonal music. Can you assign a single number to something that can vary on so many axes?
- A compression technique might be closer. Plus it would just spit out a number…
- I didn’t get that far, though.

Summary

- Factor oracle is a good stupid thing for music generation
- However, it's still stupid. You can't make a silk purse out of a sow's ear.
- It would be a good candidate for this experiment, but
- You need a good way to measure informatia, which, if you think about it, is actually a very hard problem

References

- Biederman, I., & Vessel, E. A. (2006). Perceptual Pleasure and the Brain. American Scientist, 94, 249-255
- Yue, Vessel, Biederman (2007). The neural basis of scene preferences. NeuroReport v.18 n. 6