Generative Memory

by Bret Fontecchio

Memory in Humans

- "Memory is like fiction; or else it’s fiction that’s like memory." - Haruki Murakami
- Memory in humans could be perfect. The brain probably doesn’t need to use lossy memory to save space.
- Many memories in humans remain subconscious, or they change substantially as they are recalled.
- For example, if I get into a traffic incident, I might remember it being a little less my fault.
- Some people even remember distinctly that the other guy said something that in fact they said
- Human Memory is by no means an exact regurgitation of the events.

Memory in Humans

- Prior assumptions play heavily into human perception

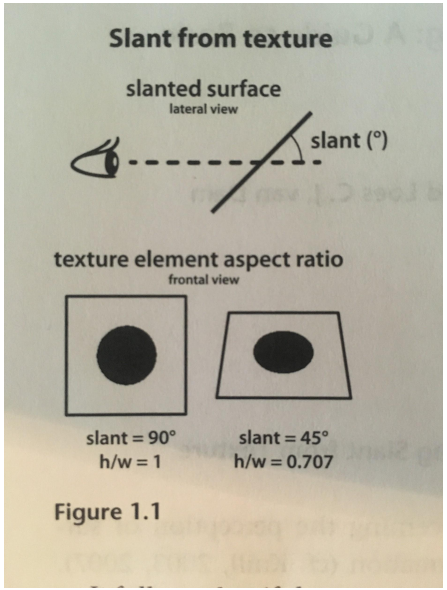

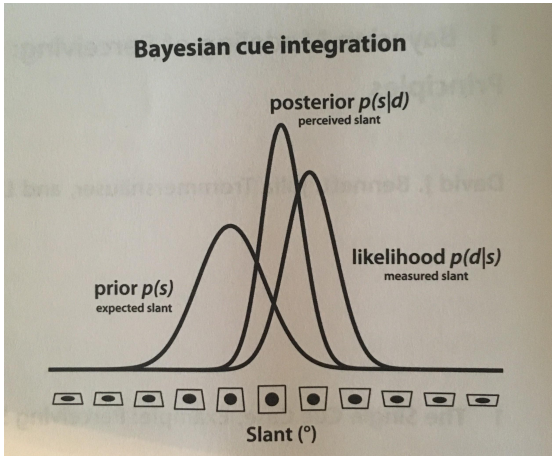

Bennett, David J., and Christopher S. Hill. Sensory Integration and the Unity of Consciousness. Print.

Bayesian Priors

So if I view a slanted object and use binocular vision to gauge how deeply slanted it is, this will be convoluted with my assumptions about slant given a Bayesian prior. I can see the outline of the object and this implies a certain slant if the lines are indeed straight when viewed straight on. So this is an example of Bayesian Priors being used by the brain.

Deep Beliefs (in humans)

- Human beliefs can be thought of as having layers
- I might believe I said something, but at a deeper layer I believe that it was spoken in my voice. A deeper layer contains beliefs about the tone and frequency.

Deep Beliefs (in machine learning)

- “Deep” refers to Artificial Neural Network layers
- More than just a buzzword
- Multilayer models are exponentially more efficient than single layer models, although in theory, given infinite time, you only need one layer (Bengio 2007)

Guiding Principles

- Take inspiration from Biology
- If it worked for millions of years it is successful
- Copying Biology is not always the right approach, but in the right context in can be.
- AI has its own goals so “that’s not how neurons/the brain work” doesn’t discredit the system although biology is the inspiration.
- Ultimately the usefulness of the system will be the rubric upon which it is judged.

Guiding Vision

I'd like to eventually get to a point where I can take audio data from someone's speech along with the phonemes connected to it, and create a biased but believable rendering from memory. So I can take prior assumptions about what they probably said, memory of their actual words and tone, and other information like what a Boston accent usually sounds like, and generate a plausible account.

Proof of Concept System

- A solid proof-of-concept system might use words only, instead of focusing on complex tasks like remembering audio and facial data
- One idea would be a system that reconstructs sentences using prior assumptions.
- Word frequency tables have been used for sentence reconstruction
- Other, more interesting “beliefs” can we worked into sentence reconstruction
- A Deep Belief Network which reconstructs sentences based on higher level, abstract beliefs about the speaker and the subject matter and biases an account would be an interesting step toward a generative AI memory

Implementation

How it snowballs:

Multilayer perceptron
  Predict y given X
Autoencoder
  Predict x given X
Autoencoder and autodecoder
  Recognition and generation
Denoising Autoencoder
  Use the training algorithm to clean data
Variational Autoencoder
  Use Bayes' Rule to bias the training
Backpropagated RBM Sub Layers
  Deep Belief Network

Implementation Cont

- Deep Belief Networks
- Layers of Restricted Boltzmann Machines
- Contrastive Divergence at each layer

Implementation Cont (2)

The training algorithms at each layer are the key here.

Relevant Software Packages

- Theano
- Anacondas
  - Numpy
  - Scipy
  - Pandas
  - Etc.
- Scikit

Overview of the Literature

- Hinton (possibly the biggest name for DBN)
  - A Fast Learning Algorithm for Deep Belief Networks
    - https://www.cs.toronto.edu/~hinton/absps/fastnc.pdf
  - Reducing the Dimensionality of Data with Neural Networks
    - http://science.sciencemag.org/content/313/5786/504.abstract
  - A summary paper
    - http://www.cs.toronto.edu/~hinton/absps/tics.pdf
- Bengio et al
  - Greedy Layer Wise Training of Deep Networks
    - http://www.iro.umontreal.ca/~lisa/publications2/index.php/attachments/single/24
  - Representational Power of Restricted Boltzmann Machines and Deep Belief Networks
    - http://www.iro.umontreal.ca/~lisa/publications2/index.php/attachments/single/22
  - Learning Deep Architectures for AI
    - http://www.iro.umontreal.ca/~lisa/pointeurs/TR1312.pdf
- Nature Deep Review
  - http://www.cs.toronto.edu/~hinton/absps/NatureDeepReview.pdf

"Memory is like fiction; or else it’s fiction that’s like memory." - Haruki Murakami

Other Directions

- 2vec systems
  - word2vec, face2vec, etc
  - Bring everything to the same lower dimensionality space and do transformations
- Stacking autoencoders
  - Instead of Restricted Boltzmann Machines
- Hopfield Networks
- LSTM’s