

Machine Learning

3 Trends That Matter and 10 That Don't

Glen Anderson
Discussion topics

Quick backgrounder on Glen

- VLSI designer 1990-1994 (Sun UltraSPARC V9 64-bit MMU)
- Electronic Design Automation companies 1994-2000
 - Ambit with Rajeev
- Some angel investing, VC advising, consulting 2000-2006
- Google 2006-present
 - 7 years in Platforms Networking (cluster fabrics, congestion)
 - 4 years in Machine Learning (Sibyl, then Brain team)
- Cryptocurrencies
- Glen's proposed "Rule of Mom" for angel investing
 - Cafepress, InBrace

Basics

Regression, classification, sequence, clustering, generative, RL

Dense vs. sparse (categorical) inputs

Supervised vs. unsupervised

Classification

<https://aiexperiments.withgoogle.com/>

Representations

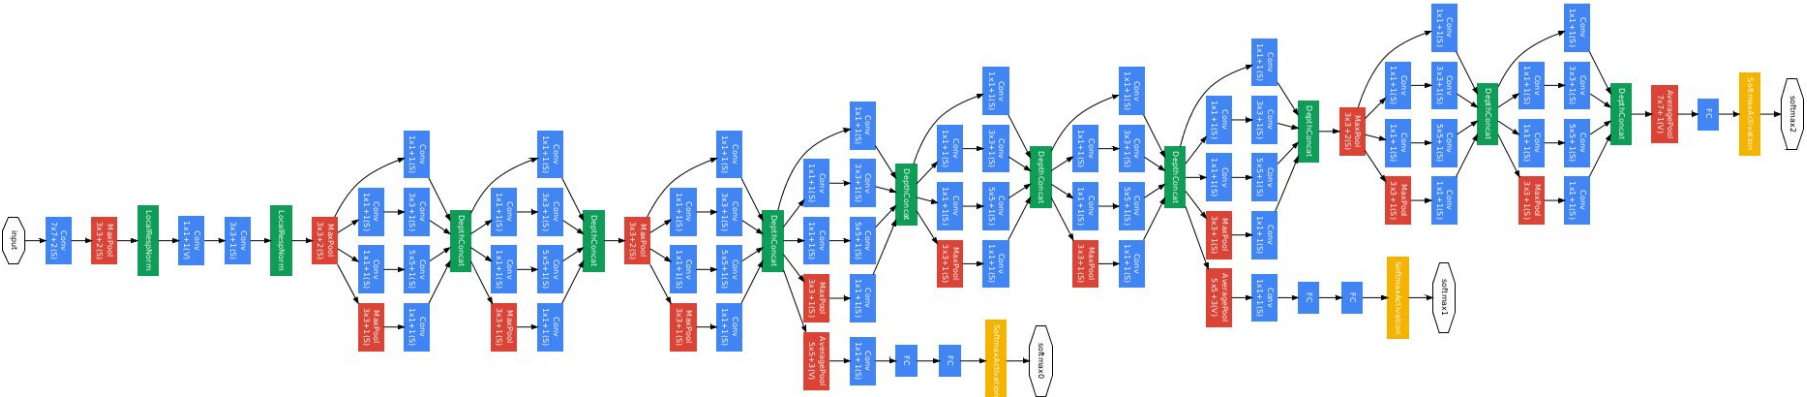
<https://aiexperiments.withgoogle.com/visualizing-high-dimensional-space>

<http://projector.tensorflow.org/>

Understanding how

<http://playground.tensorflow.org>

Deeper and deeper (for image/video)



GoogLeNet model architecture

Combined visual and semantic embeddings

Nearest Images



- day + night =



- flying + sailing =



- bowl + box =



- box + bowl =



Nearest images



- dog + cat =



- cat + dog =



- plane + bird =



- man + woman =



A person riding a motorcycle on a dirt road.



Two dogs play in the grass.



A skateboarder does a trick on a ramp.



A dog is jumping to catch a frisbee.




A group of young people playing a game of frisbee.



Two hockey players are fighting over the puck.



A little girl in a pink hat is blowing bubbles.



A refrigerator filled with lots of food and drinks.



A herd of elephants walking across a dry grass field.



A close up of a cat laying on a couch.



A red motorcycle parked on the side of the road.



A yellow school bus parked in a parking lot.



Describes without errors

Describes with minor errors

Somewhat related to the image

Unrelated to the image

Attentional models



A woman is throwing a **frisbee** in a park.



A **dog** is standing on a hardwood floor.



A **stop** sign is on a road with a mountain in the background



A little **girl** sitting on a bed with a teddy bear.



A group of **people** sitting on a boat in the water.



A giraffe standing in a forest with **trees** in the background.


Running networks in reverse



Hartebeest



Measuring Cup



Ant



Starfish



Anemone Fish



Banana



Parachute



Screw

<https://deepdreamgenerator.com/>
<https://deepdreamgenerator.com/gallery>

<https://deepdreamgenerator.com/ddream/d024edebe6>




<https://deepdreamgenerator.com/ddream/e1bc24baea>




<https://deepdreamgenerator.com/ddream/c7e91a713e>




Adversarial inputs



dog



+noise



ostrich

GANs (Generative Adversarial Networks)


<https://affinelayer.com/pixsrv/>

Memory consolidation experiment

(credit for the following to DeepMind)

Hippocampus sits at the center of the brain--integrates the sensory information.

It stores short-term memories (observed relationships between data)



Sleep consolidation

During sleep, the "training data" from your day's experience is replayed from the Hippocampus to the cortex, many times

They can measure a rat "running a maze" during sleep

The cortex efficiently encodes and stores this information

Slow oscillation
(0.5-1.0Hz;
neocortex)




← 1.0s →

Spindle
(11-15Hz;
thalamus)

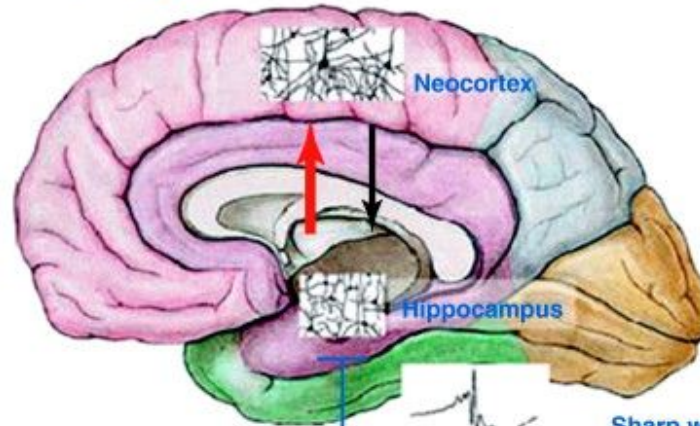


← 0.5s →

Sharp wave ripple
(100-300HZ;
hippocampus)



← 0.2s →



Sharp wave-ripple
memory replay

Acetylcholine,
cortisol, etc.

Sleep consolidation experiment (1)

Try this on your family and friends:


1. Print out 7-9 images of simple objects onto paper, cut them out
2. Order them (arbitrarily) but don't show your test subject the images or tell them you ordered them



Sleep consolidation experiment (2)

3. In the evening, do a test:

3a. Quiz them with 20 adjacent pairs, "Who wins?" you ask. At first they have no idea what you are talking about. You record right/wrong and tell them the answer. They learn! By pair 10 they should get them all right.




Sleep consolidation experiment (3)

Remember: at this point they don't know the underlying theme

All they see are pairs of objects and they start to memorize the pattern:

"flower always beats shoe"

"bread always beats airplane"




Sleep consolidation experiment (4)

4. In the same evening session

4a. Quiz them with 10 random pairs. Record their answers, right/wrong. (Probably 50% wrong)

-> "dog, bread... no idea!"


Don't tell them these answers!



Sleep consolidation experiment (5)

5. The next morning:

5a. Quiz them with 10 new random pairs. Record their answers, right/wrong.




Sleep consolidation experiment (6)

If this works correctly:

Not only do they get all the answers correct in the morning...

they insist you showed them those exact pairs last night! The correct answer is that obvious them.



Sleep consolidation experiment (6)

"Of course flower beats house"

The cortex has organized the information efficiently, maybe in a 1-dimension embedding!

See also e.g.

<http://psychology.msu.edu/SleepLab/MemoryTest.aspx>



Theory of humor

If a brain benefits from organizing information efficiently, what's a good motivation for evolution to encourage it?

What if a small pleasure signal was given when a long series of jumbled information finally made sense with a punch line, allowing the storage to be consolidated?

Or the satisfaction of understanding a complex topic?


TREND 1 - End-to-end learning

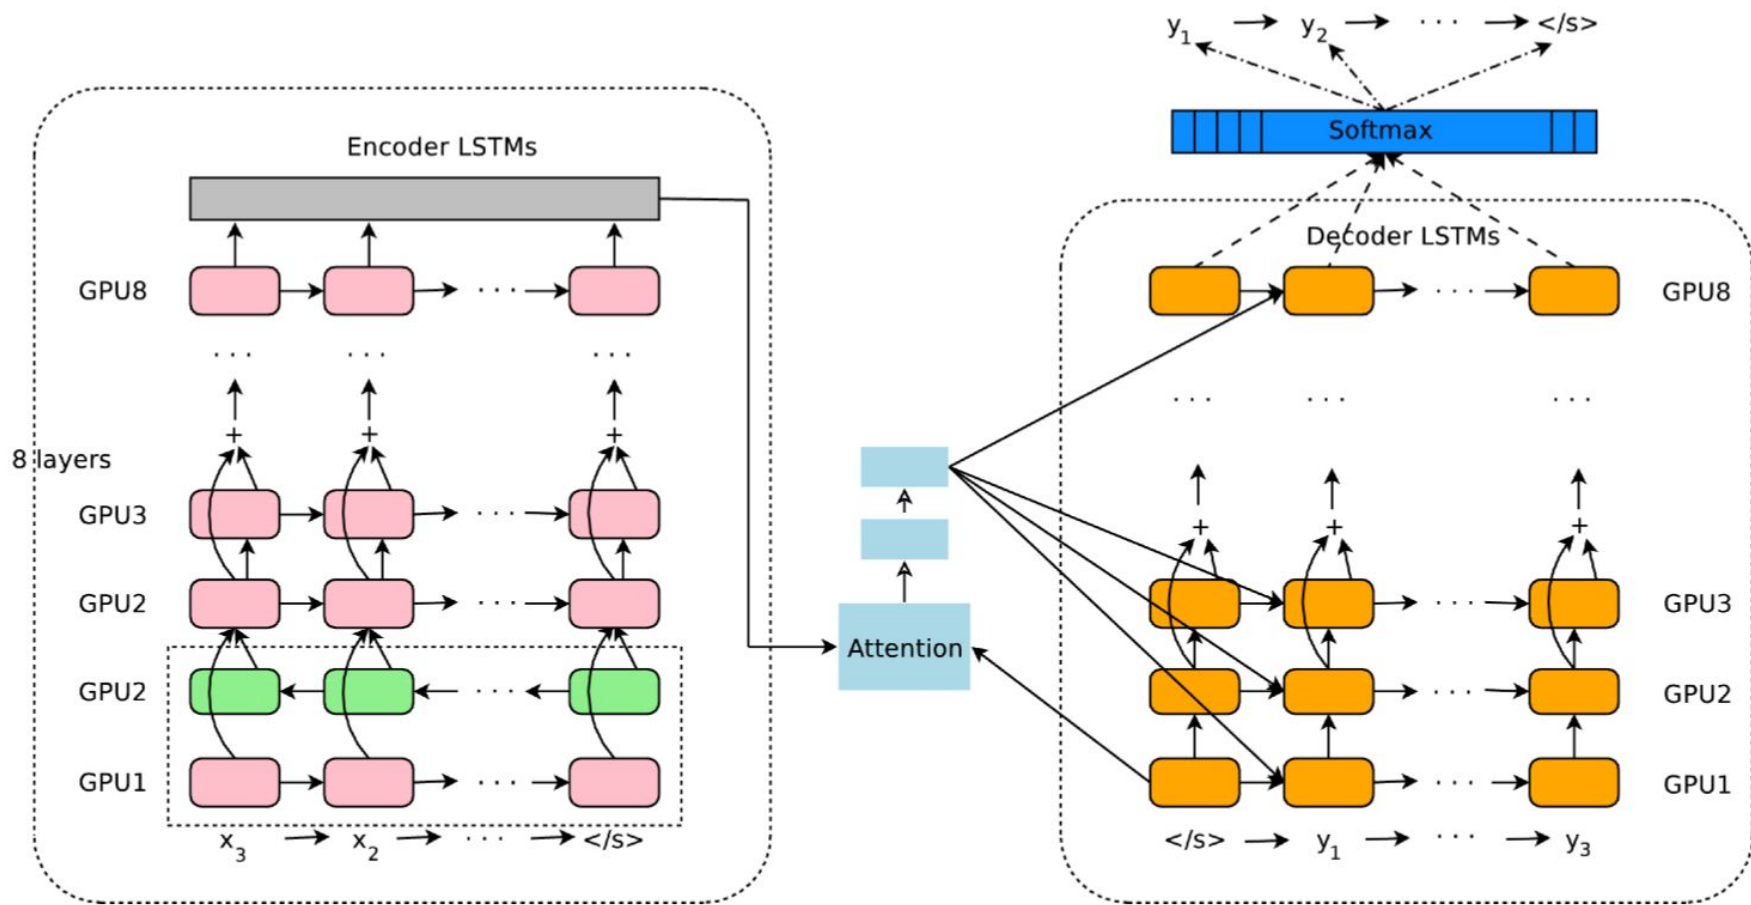
- Language translation
- Speech recognition
- Robotics
- ... many more coming



<https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>

Google Neural Machine Translation





Zero shot learning

- High-level representations allow surprising applications
 - Learn to translate A→B
 - Learn to translate C→D
 - Ask the system to translate A→D

TREND 2 - Data

TREND 3 - The commoditization of algorithms

- Diminishing gains at any one task
 - benchmarks
- TensorFlow for published papers
 - great for humanity
 - maybe not so great for technical entrepreneurs

Trends that might not matter (*)

(*) to an entrepreneur; to an engineer they may be interesting

- MLaaS
 - offers nothing new
- Custom accelerators
 - quirky
- Fairness in ML
 - no wins
- Privacy
 - bogged down
- Medical
 - too many barriers, slow
- Massive networks
 - aggregation
- Mixture of experts
 - aggregation
- Custom loss functions
 - too incremental
- Parameter tuning
 - too incremental
- Massively sparse
 - maybe