THE FUTURE OF AUTOMATED HEALTHCARE

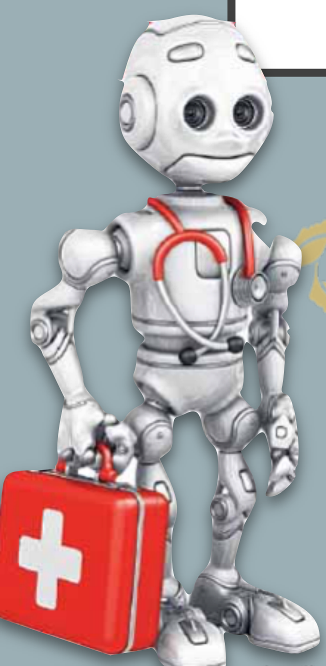

Frank Rudzicz

UHN Toronto Rehabilitation Institute

University of Toronto

Winterlight

Vector Institute

THE PRESENT OF AUTOMATED NEWS

- 61% of millennials get their news from Facebook (Pew Research).
- Facebook wants to **maximize your 'engagement'** with its news posts.
  - "Engagement is good" – you get to interact with the news that matters to you (and Facebook can use your preferences to sell you to advertisers).
- Facebook's AI algorithms **sort** potential news to show you; something *like*:

$$P(engagement|\;) = \lambda_1 P(Like|\;) + \lambda_2 P(comment|\;) + \cdots$$

- Machine learning optimizes weights λ_i and the probabilities themselves.

YourNewsWire.com

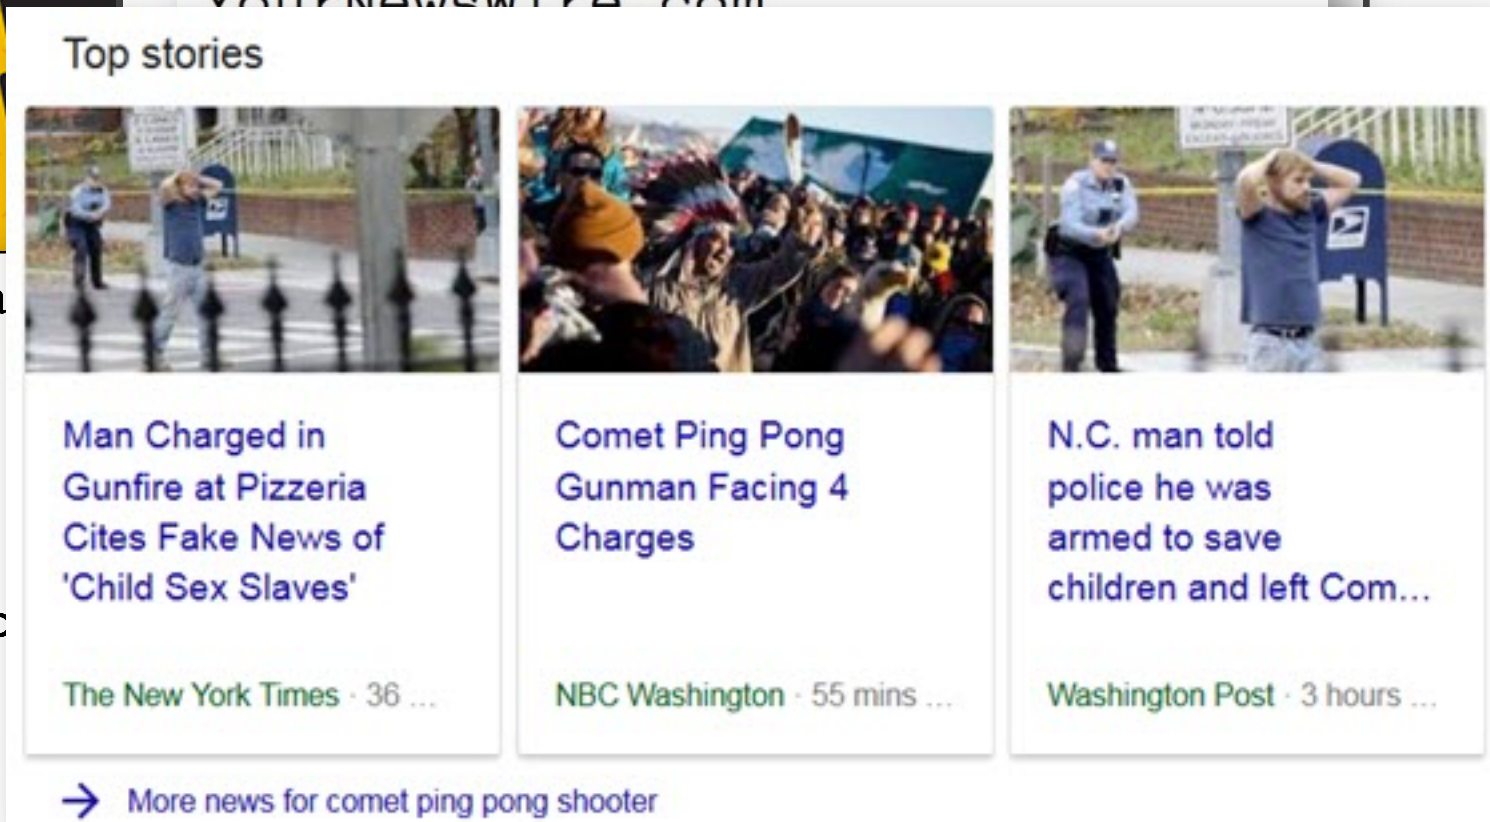

all

- Fa
- Yo ike
  ent,
  as),
  c) …
- There's a finger on the scale. A feedback loop.

THE PRESENT OF AGITATED HUMANS

- The risk is not disobedient artificial intelligence, but instead *obedient* artificial intelligence mixed with lazily-defined objectives.
- Social media shows us that simple algorithms designed for innocuous purposes *can* have disastrous unintended consequences, if mixed with hidden human biases.

"First we shape our tools, and thereafter our tools shape us."
-Marshall McLuhan

- What simple algorithms are being designed in healthcare?

THE FUTURE OF AUTOMATED HEALTHCARE

- What are some **current trends** in AI for healthcare?

- What **technical risks** exist? How can we patch them?
- What **regulatory hurdles** exist? How can (should?) we surmount (circumvent?) them?
  - What changes are likely coming for our **society**?

CURRENT TRENDS IN MACHINE LEARNING

TRENDING NOW

1. Deep neural networks (of course)
2. Big Data (with cells interlinked within cells interlinked)
3. Recurrent neural networks for temporal, dynamic data
4. Reinforcement learning
5. Active learning
6. Telehealth and remote monitoring
7. Causal, explainable models

1. DEEPER, DEEPER STILL

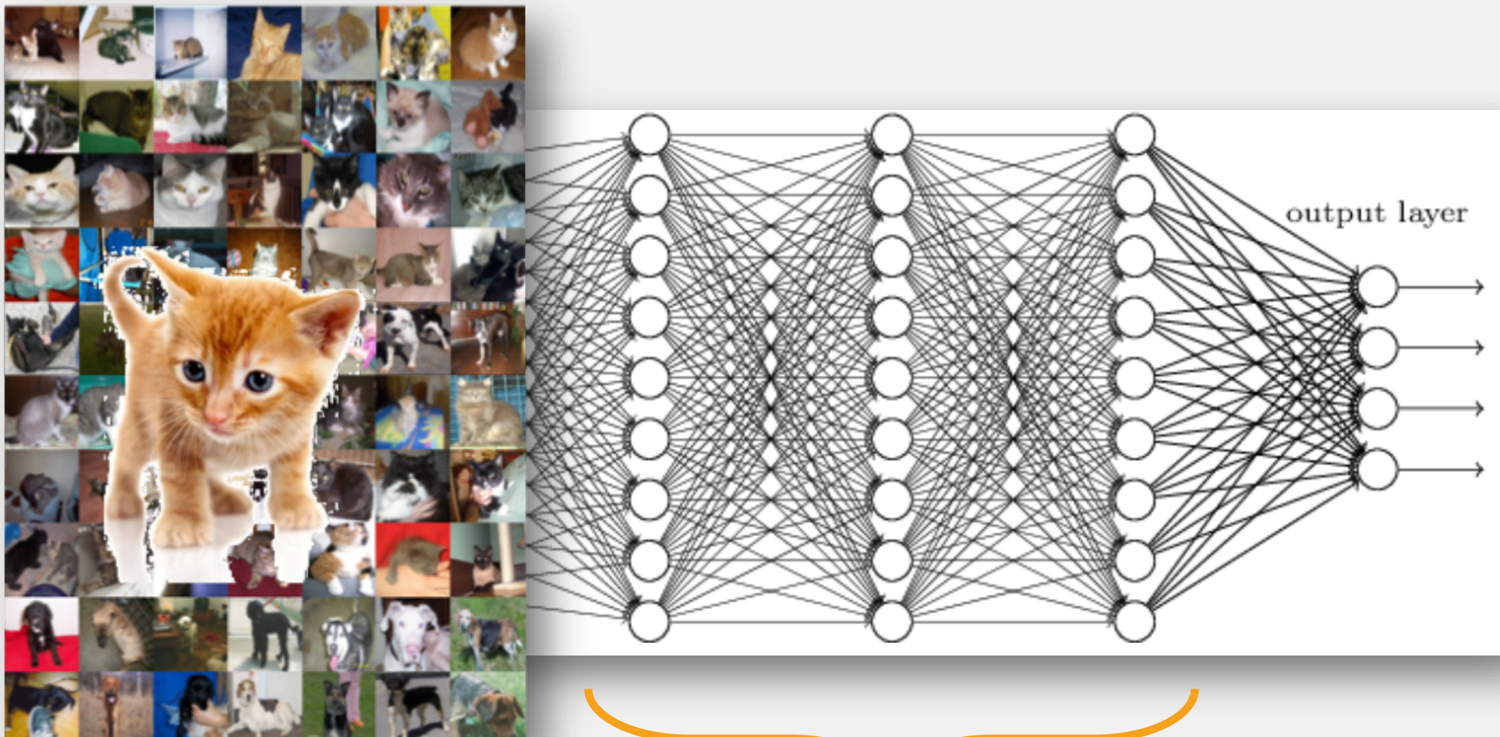

It's a freaking cat.

'hidden' representations are learned here

1. DEEPER, DEEPER STILL

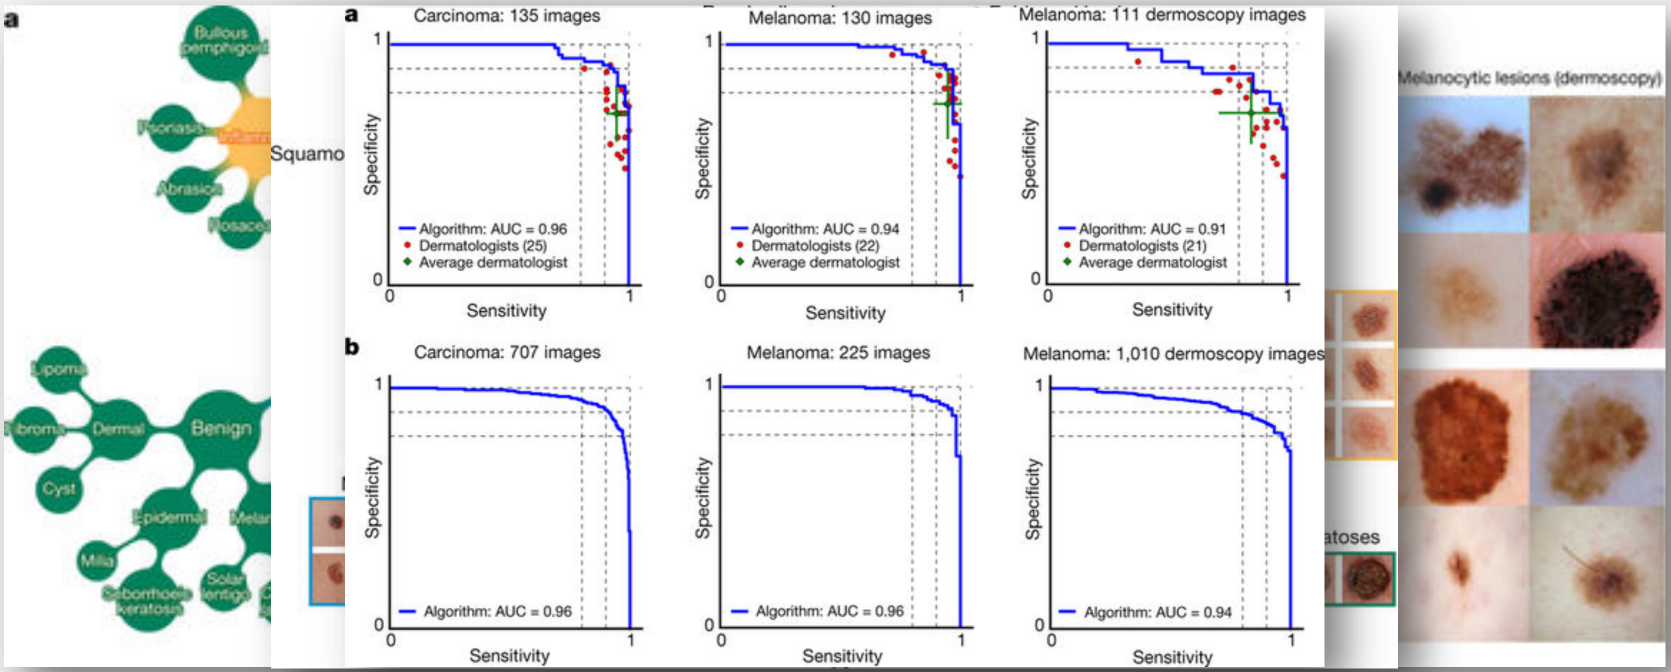

Trained with 129,450 clinical images
Tested against 21 certified dermatologists.

Esteva A, Kuprel B, Novoa RA, *et al.* (2017) Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**:115-118

2. BIG (INTERLINKED) DATA

- Modern deep learning *usually* requires a **large amount** of data to work.
- Fortunately, when the data set is very big, the need to **preprocess** it decreases.
- Supervised training needs **labels**; often in healthcare the labels come from **validated assessments** recorded in parallel.
- Sometimes we use unsupervised training, e.g., cluster analysis.
- The real promise of Big Health Data is in the interconnections that exist **between** Big Data sets.

2. BIG (INTERLINKED) DATA

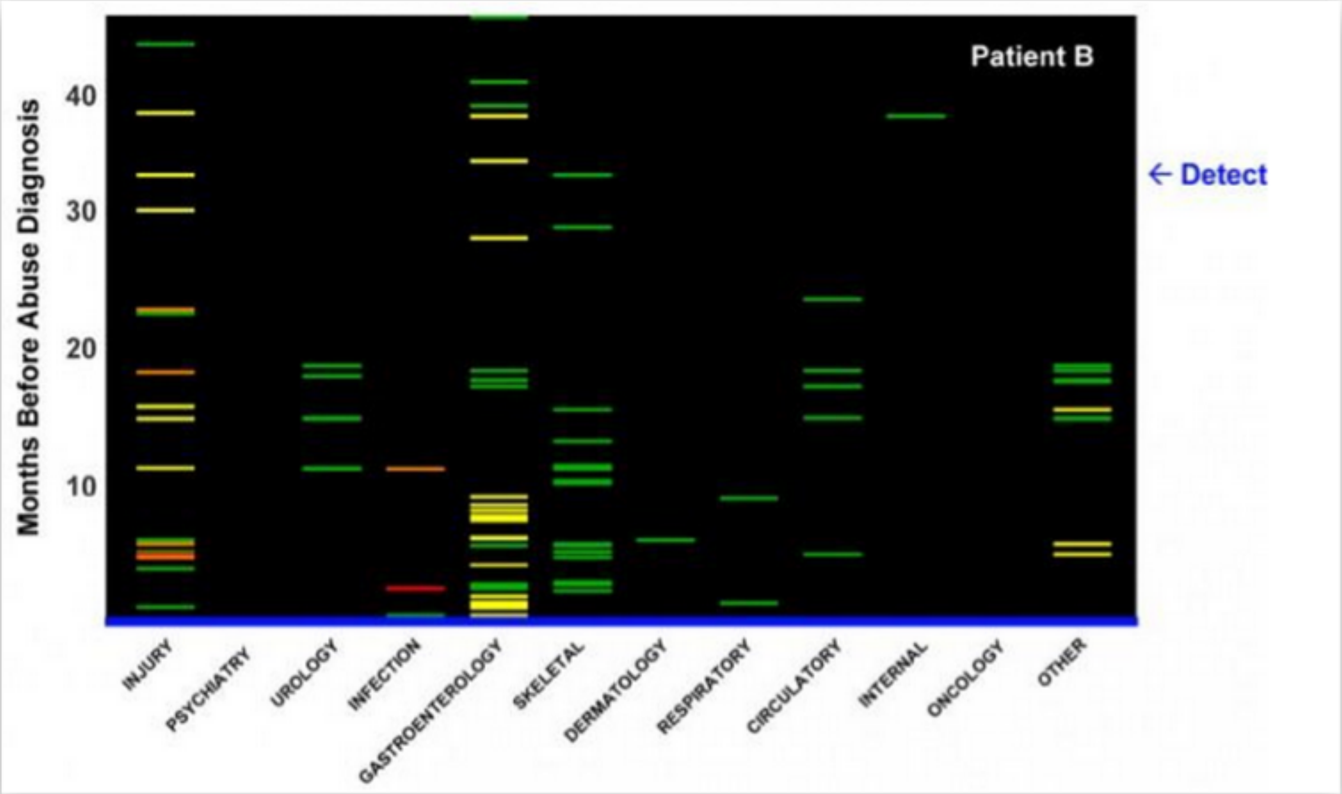

Reis, B.Y., Kohane, I. S., & Mandl, K. D. (2009). Longitudinal histories as predictors of future diagnoses of domestic abuse: modelling study. *BMJ (Clinical Research Ed.)*, **339**, b3677.

3. TIME SEQUENCES AND RECURRENT DYNAMICS

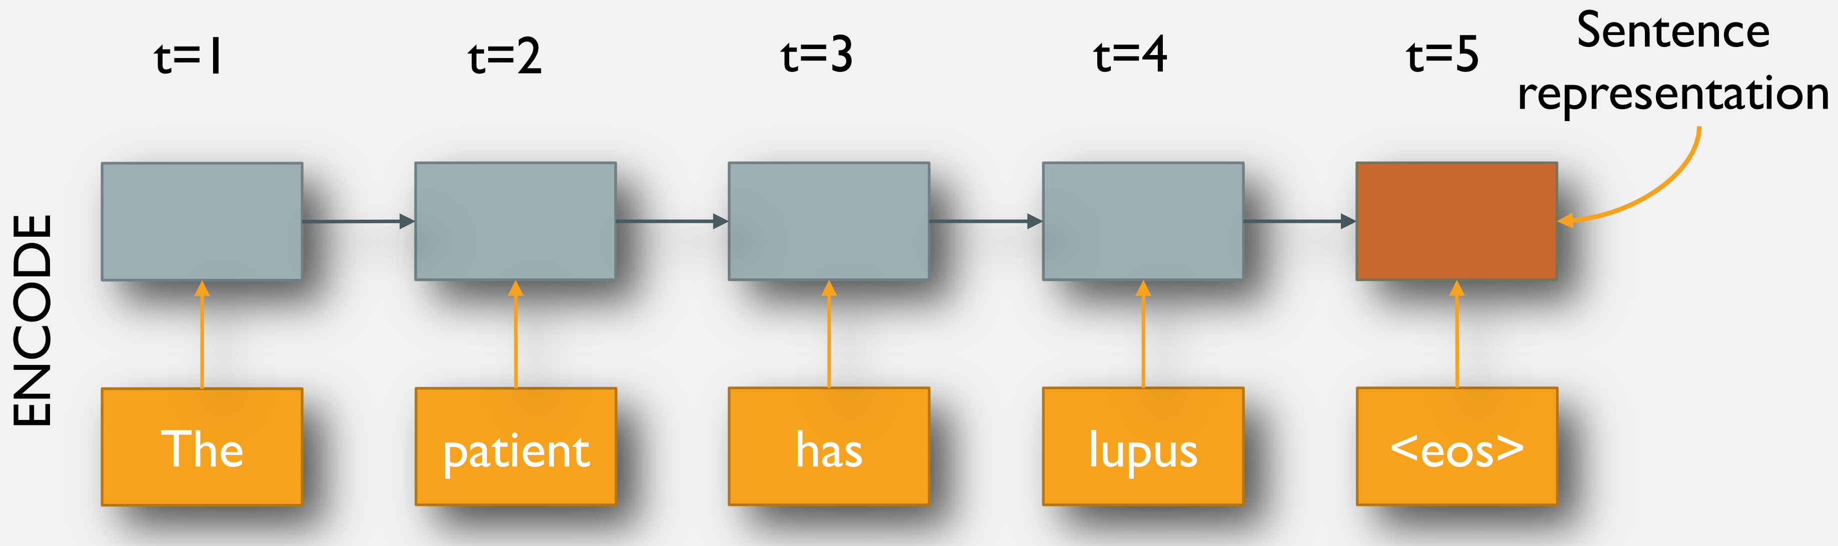

- Unfortunately, for very long sequences, this pseudo-Markov process quickly forgets earlier information.
  - E.g., this network might have almost no idea what to predict after "*I grew up in France so I'm pretty good at speaking ….*"

3. TIME SEQUENCES AND RECURRENT DYNAMICS

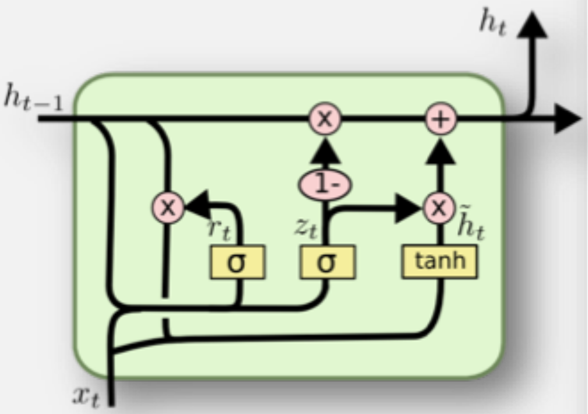

Long short-term memory networks

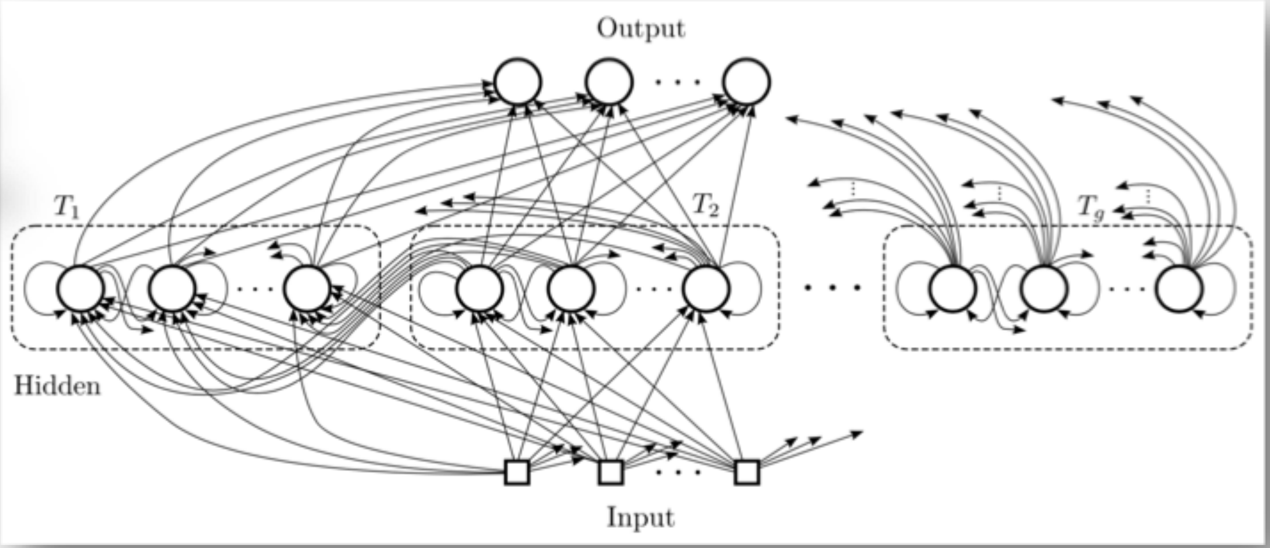

Clockwork recurrent networks

3. TIME SEQUENCES AND RECURRENT DYNAMICS

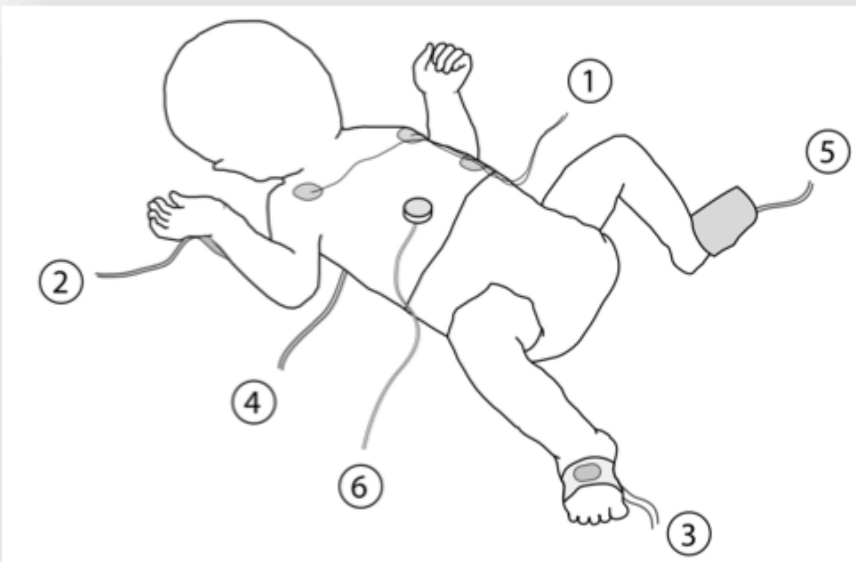

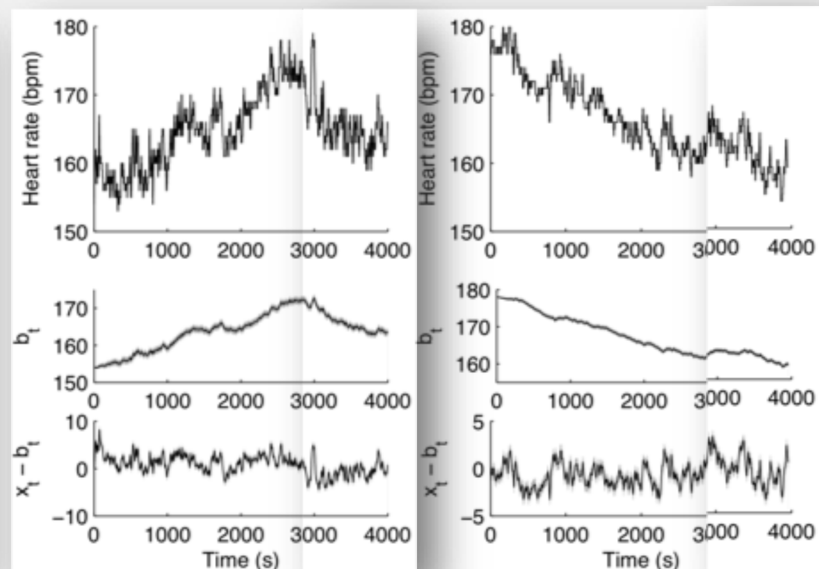

- Long-term trends are observable over milliseconds and over years.

Quinn JA, Williams CKI, Mcintosh N (2009) Applied to Physiological Condition Monitoring. 2009; IEEE TPAMI **31**:1537–51.

Lipton ZC, Kale DC, Elkan C, *et al* (2015) Learning to Diagnose with LSTM Recurrent Neural Networks. 2015, In *Proceedings of ICLR*: 1–18.

4. REINFORCEMENT LEARNING

- **Reinforcement learning** was inspired somewhat by behaviourist psychology.
  - Systems learn 'online' in the real-world (theoretically) by taking imperfect **observations**, inferring the unseen **state** of the world, and then taking an **action**.
  - Actions are chosen to maximize an expected **reward**, or to minimize an expected **cost**.
    - These rewards and costs are usually supplied by humans.
    - In order to learn, some **exploration** is necessary.

4. REINFORCEMENT LEARNING

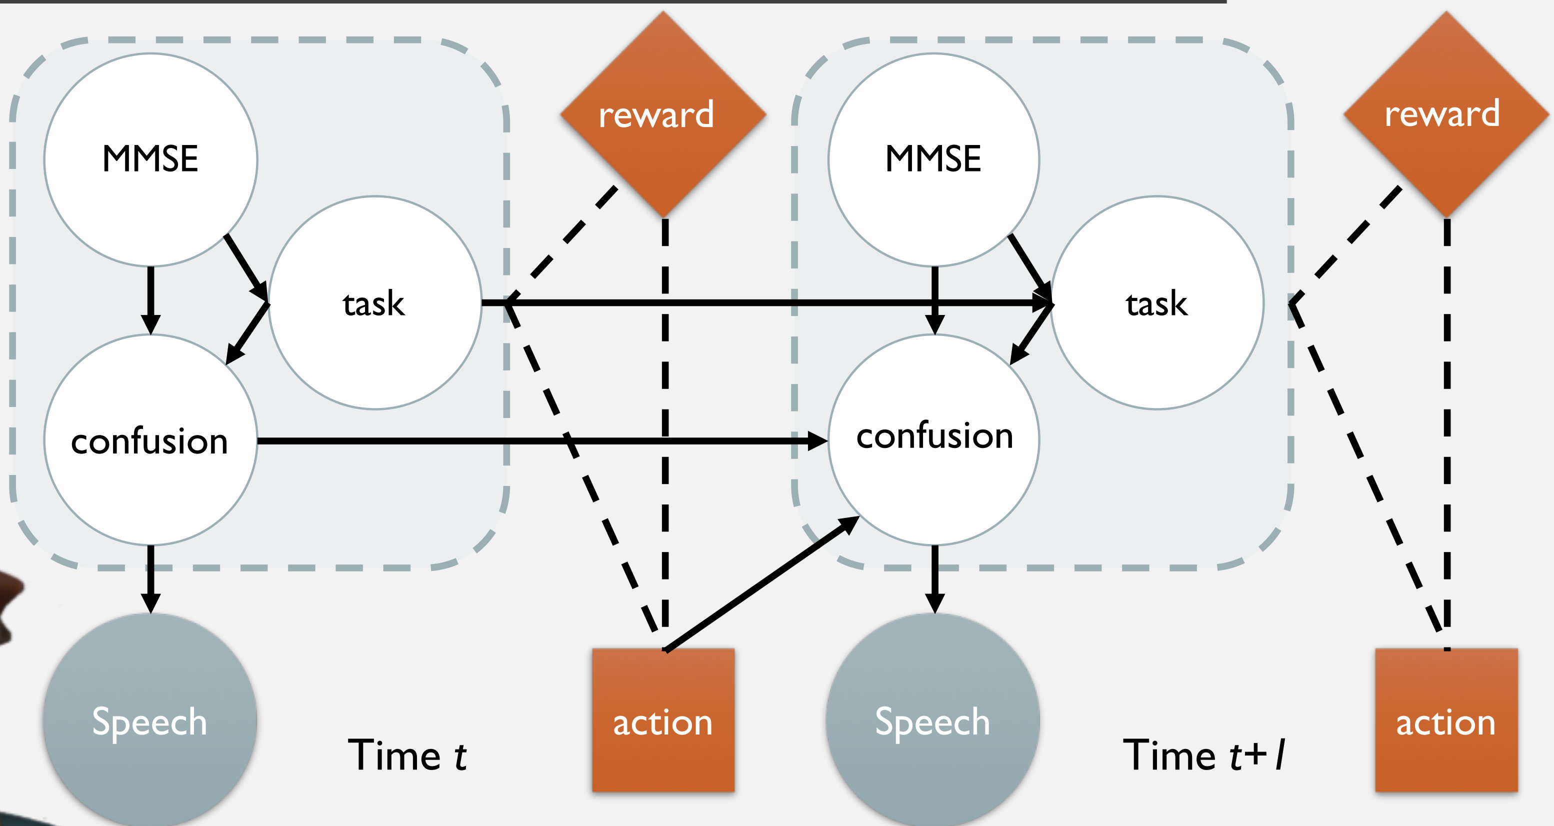

Be careful about reward hacking!

4. REINFORCEMENT LEARNING

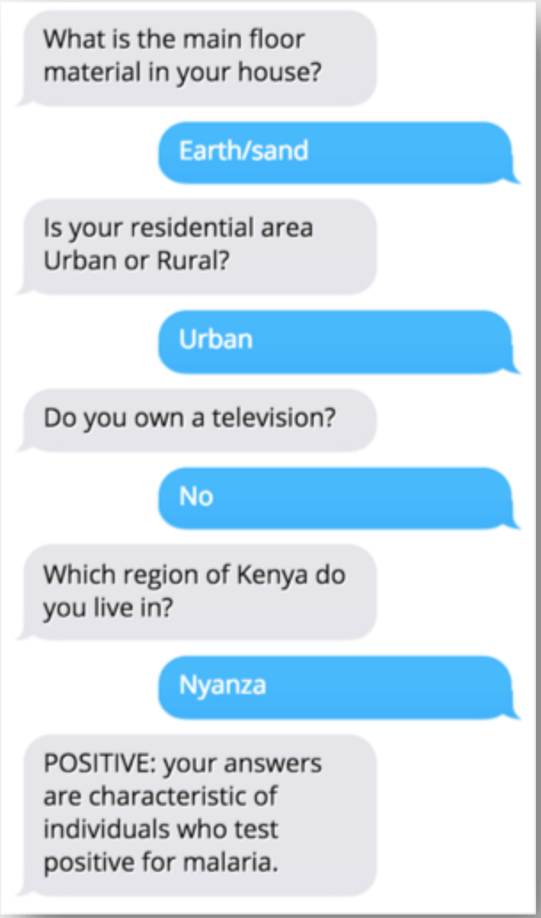

- Any interaction with a simulated human doctor will likely use **reinforcement learning** to choose what questions to ask, what labs to order, what interventions to prescribe…
  - Rewards and costs are usually supplied by humans.
  - In order to learn, some **exploration** is necessary.

Rajpurkar *et al* (2017) Malaria Likelihood Prediction By Effectively Surveying Households Using Deep Reinforcement Learning. *ML4H*.

5. ACTIVE LEARNING

National Institutes of Health (NIH) grants-supported research

ARTIFICIAL INTELLIGENCE FOR COMPUTATIONAL PATHOLOGY

Image interpretation plays a central role in the pathologic diagnosis of cancer. Since the late 19th century, the primary tool used by pathologists to make definitive cancer diagnoses is the microscope. Pathologists diagnose cancer by manually examining stained sections of cancer tissues to determine the cancer subtype. Pathologic diagnosis using conventional methods is labor-intensive with poor reproducibility and quality concerns. New approaches use fundamental AI research to build tools to make pathologic analysis more efficient, accurate, and predictive. In the 2016 Camelyon Grand Challenge for metastatic cancer detection,[69] the top-performing entry in the competition was an AI-based computational system that achieved an error rate of 7.5%.[70] A pathologist reviewing the same set of evaluation images achieved an error rate of 3.5%. Combining the predictions of the AI system with the pathologist lowered the error rate to down to 0.5%, representing an 85% reduction in error (see image).[71] This example illustrates how fundamental research in AI can drive the development of high performing computational systems that offer great potential for making pathological diagnoses more efficient and more accurate.

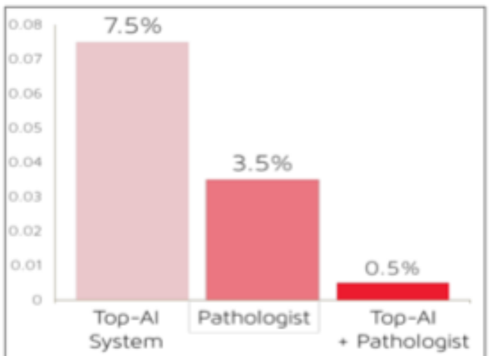

AI significantly reduces pathologist error rate in the identification of metastatic breast cancer from sentinel lymph node biopsies.

Holzinger, A. (2016). Interactive Machine Learning for Health Informatics: When do we need the human-in-the-loop? *Springer Brain Informatics*, *3*(1), in print. http://doi.org/10.1007/s40708-016-0042-6

Our Facebook scenario from before was basically a form of active learning.

6. READING YOUR MIND FROM KILOMETERS AWAY

- **Language** provides a detailed lens into human **cognition** and **sentiment**.
  - Language decline is an early hallmark of Alzheimer’s.
  - Specific brain regions serve specific linguistic functions.
  - How to measure the linguistic symptoms of cognitive decline?

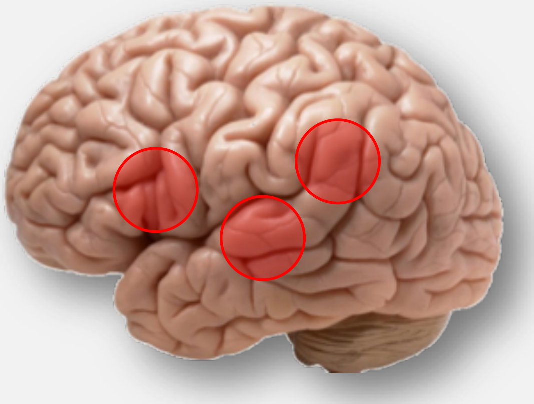

6. ASSESSING ALZHEIMER'S ON YOUR COUCH

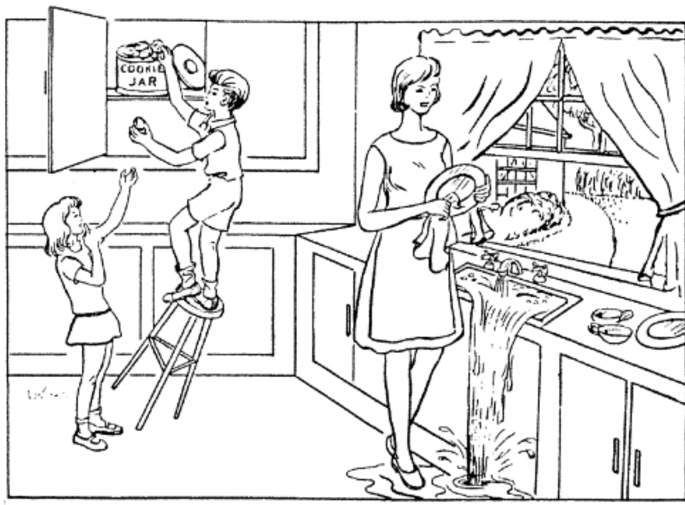

- A **picture description** task can be done in less than a minute, on the couch.

- **DementiaBank:**
  240 samples from167 people with AD,
  233 samples from 97 controls.
  - Free-form descriptions of "Cookie Theft" (incl. audio)
  - Transcribed and annotated, e.g., with filled pauses, paraphasias, and unintelligible words.
  - **Mini-mental state exam (MMSE)**

6. ASSESSING ALZHEIMER'S ON YOUR COUCH

Lexical	Frequency; Avg. word length;
Sy	;
Ac	Mean RPDE; Mean power; Pause::word ratio

WINTERLIGHT

winterlightlabs.com

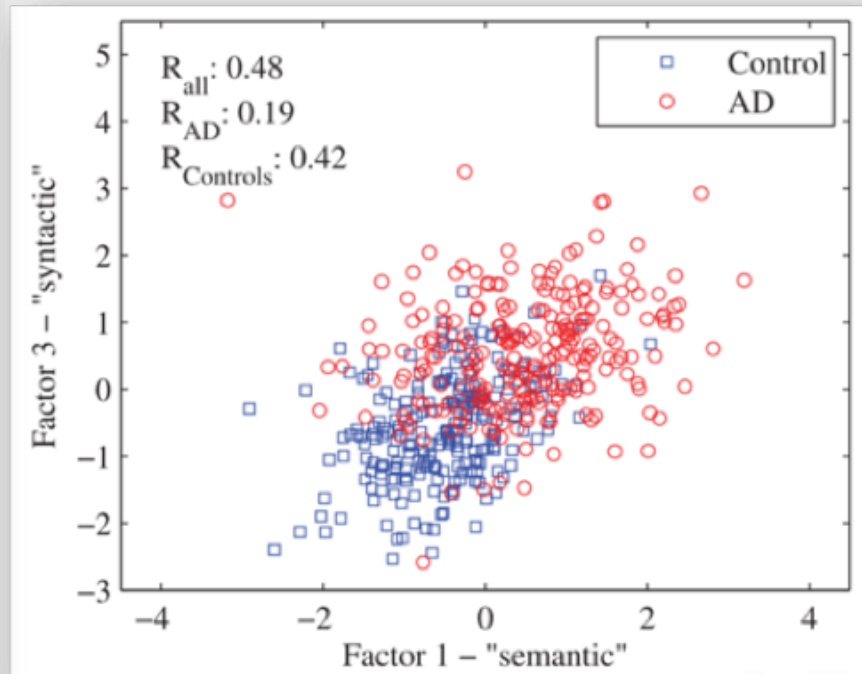

State-of-the-art accuracy: 85% - 92%

6. NEUROPSYCHIATRY ON TWITTER

- Very similar approaches can be taken for **neuro-psychiatric disorders** such as **depression**, **anxiety**.
- Traditional bag-of-words approach used dictionaries of **happy** and **sad** words, simple counts, and regression.

tweets	Hamilton Rating for Depression
best day of my life	0/50
sunny and pleasant, despite some rain	8/50
I'm glad this stupid sunny day is over	19/50

6. NEUROPSYCHIATRY WITH SIRI

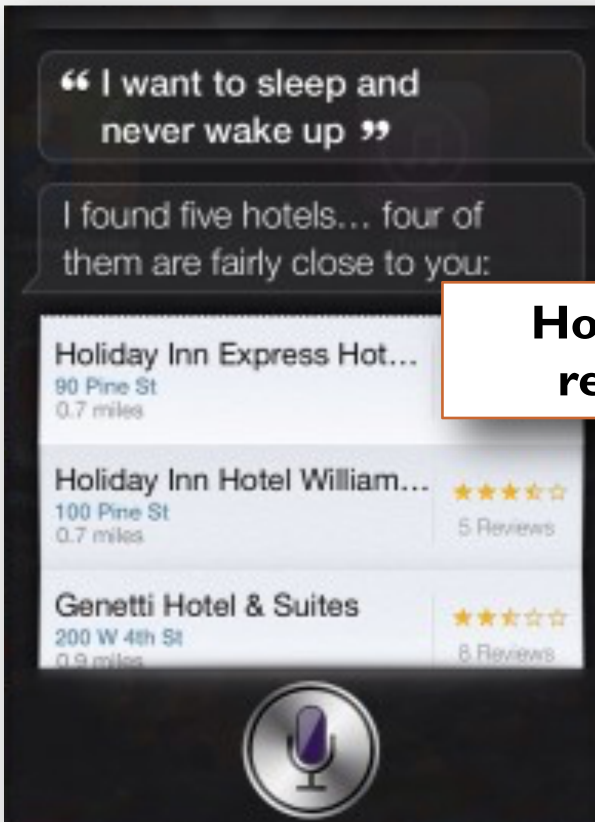

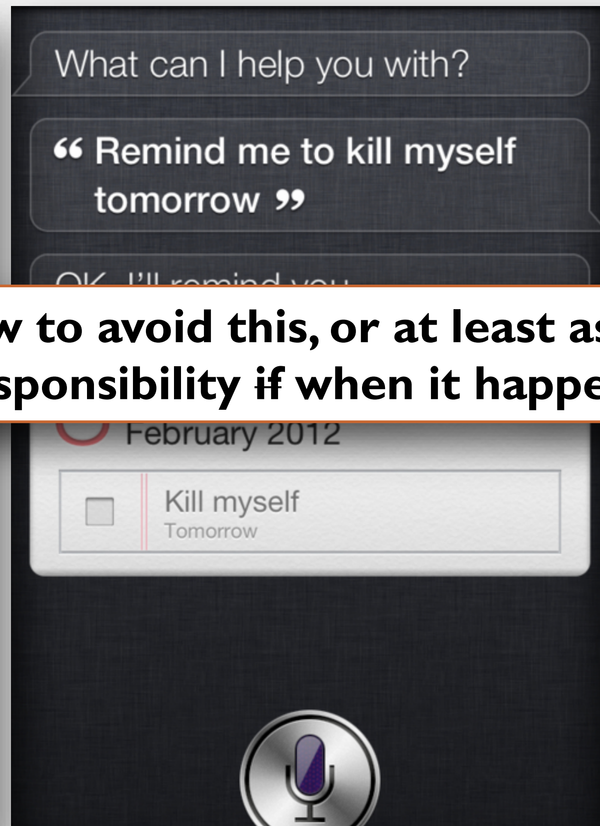

How to avoid this, or at least assign responsibility ~~if~~ when it happens?

7. CAUSAL, EXPLAINABLE MODELS

- Two criticisms of neural networks:
  - Correlation is not causation. Neural networks learn to associate input features with output categories, but there is no **abstract logic** or **interpretable reasoning** to those associations.
    - It is not logical to say a particular biomarker *causes* the system to identify a particular disease; certainly not for a new case.
  - You can't explain that. Neural networks are just matrices of $\mathbb{R}$ numbers. You can't tell *why* or *how* a neural network made a decision; you can't assign responsibility.

-0.6470	-0.2511	0.2197	-0.2381	0.7059	2.3909	0.7297	-0.4485	0.
[illegible]	0.5365	1.1141	0.9915	0.1788	0.0271	0.5102	-0.6287	-0.
[illegible]	-0.4343	0.5221	-0.1991	0.2987	-0.5563	1.2575	-0.0974	0.
-2.4764	1.4466	-0.1336	1.3179	-0.8452	0.0760	-0.9498	0.5243	1.
1.1600	0.8925	-0.3220	0.1448	-0.0111	1.3068	-0.1880	0.3389	1.
0.8969	-0.5337	-1.3633	-0.2490	0.8003	-0.0838	-0.3087	-0.4182	0.
-1.2177	-0.2749	-0.6337	-1.6915	-1.4943	1.5698	1.4739	1.5226	0.
-0.7138	-0.6191	-0.1737	2.2956	-0.3934	-0.8039	2.3219	-1.3244	1.
0.3614	0.6738	-0.3593	0.7196	-1.3263	-1.3081	-0.3592	-1.3906	0.
-1.3505	-2.3676	-1.8209	-0.4147	1.3381	-0.7390	0.4410	2.4876	0.
-1.0767	0.8864	-0.8456	0.2035	0.4567	-0.5647	-1.3996	0.3438	0.

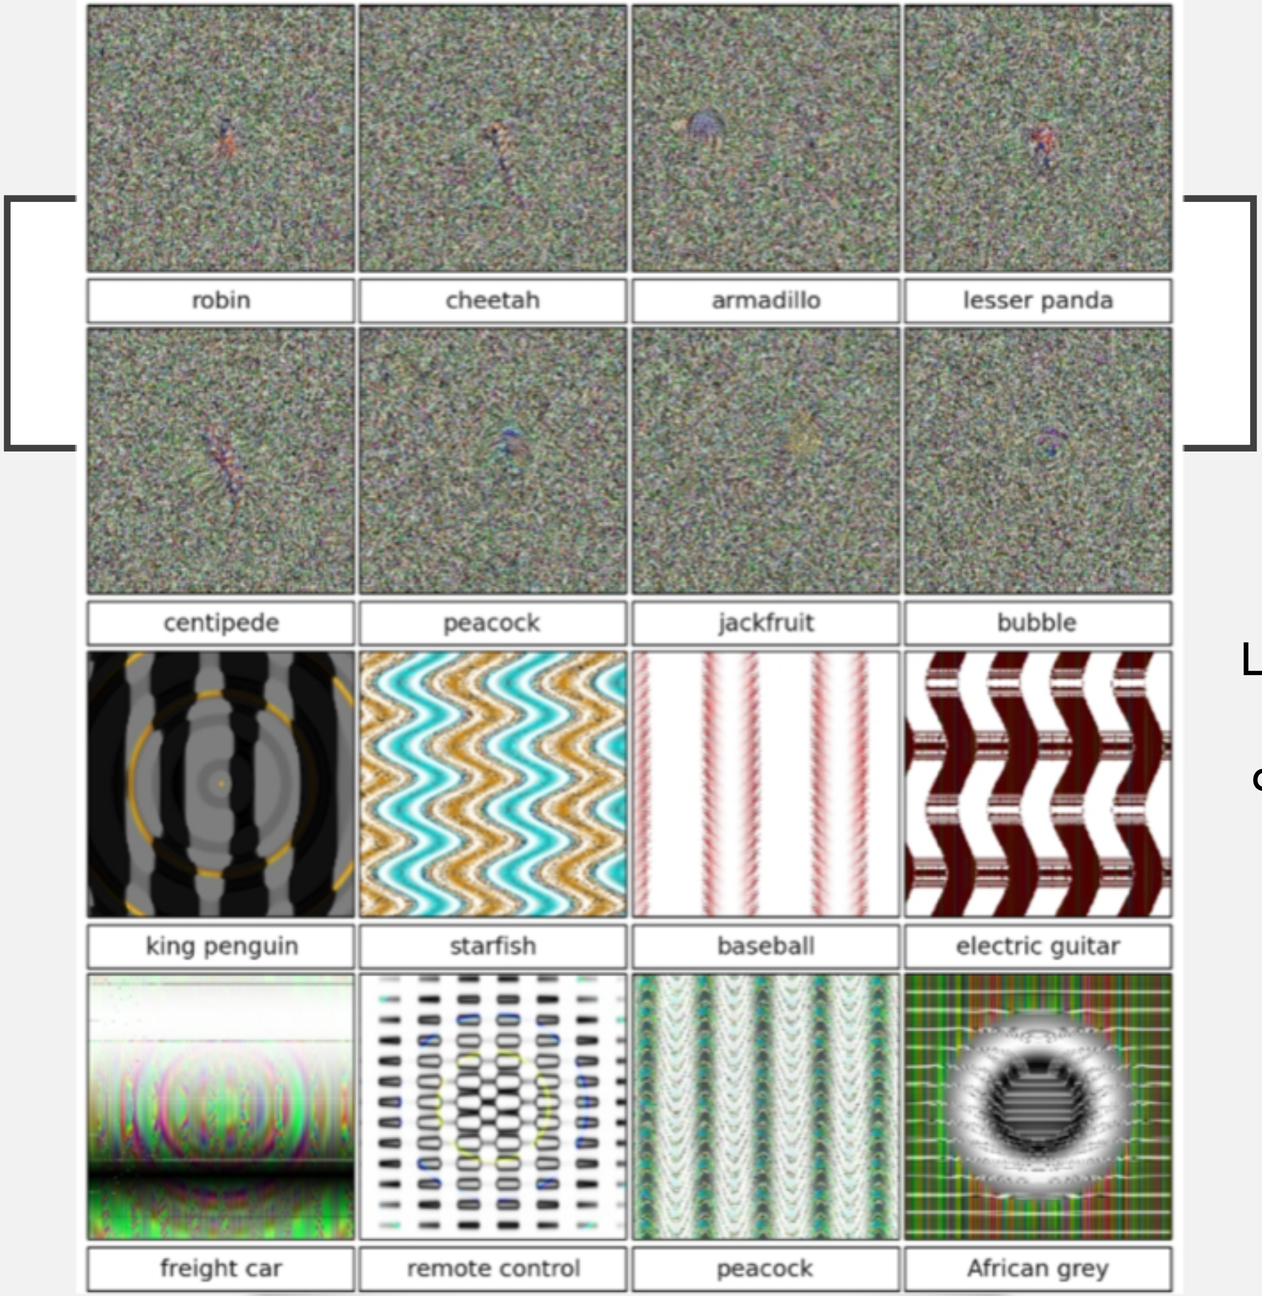

Labels with >99% confidence

Nguyen A, Yosinski J, Clune J. (2015) Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *Proc. of IEEE CVPR*. 427–36.