VANDERBILT UNIVERSITY School of Medicine

Society for Anesthesia & Sleep Medicine

Top 5 Technologies in Sleep and Anesthesia

Jon Wanderer, M.D., M. Phil
Vanderbilt University
Department of Anesthesiology

VANDERBILT UNIVERSITY School of Medicine

Artificial Intelligence and Sleep Medicine: Hope or Hype?

Jon Wanderer, M.D., M. Phil
Vanderbilt University
Department of Anesthesiology

VANDERBILT UNIVERSITY School of Medicine

Are Robot Zombies Going to Take Over Medicine?

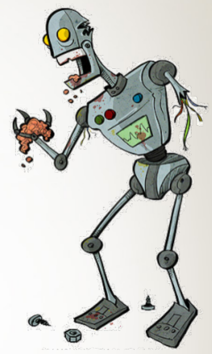

Jon Wanderer, M.D., M. Phil

Vanderbilt University

Department of Anesthesiology

VANDERBILT UNIVERSITY School of Medicine

Goals and Objectives

- Understand the context for the development of artificial intelligence in health care
- Explain how neural networks work
- Describe applications of AI in health care
  - Top 5 Technologies in Sleep and Anesthesia
- Describe limitations of the artificial intelligence

VANDERBILT UNIVERSITY School of Medicine

We live in the future

VANDERBILT UNIVERSITY School of Medicine

AI Everywhere

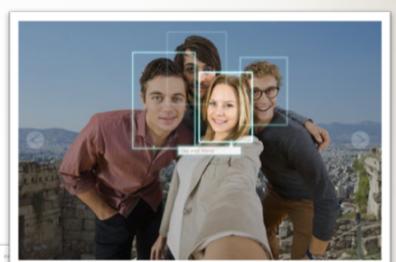

VANDERBILT UNIVERSITY School of Medicine

Moore's Law

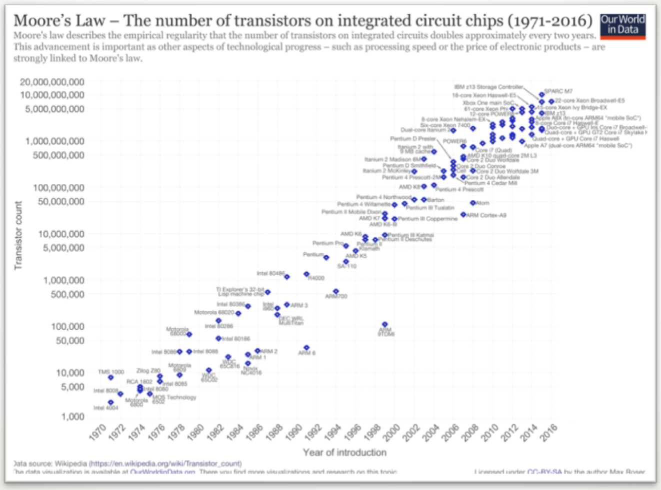

"The number of transistors that can be packed into a given unit of space will roughly double every two years."

VANDERBILT UNIVERSITY School of Medicine

Health Information Technology for Economic and Clinical Health Act (HITECH Act)

- Part of the American Recovery and Reinvestment Act of 2009 (ARRA)
- Created to motivate the implementation of electronic health records (EHR)

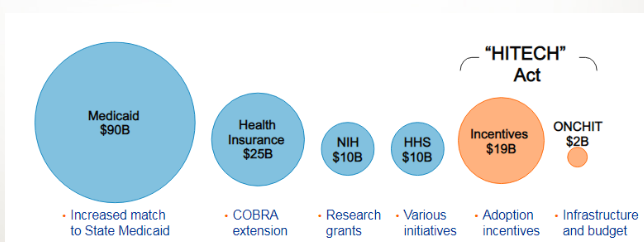

VANDERBILT U

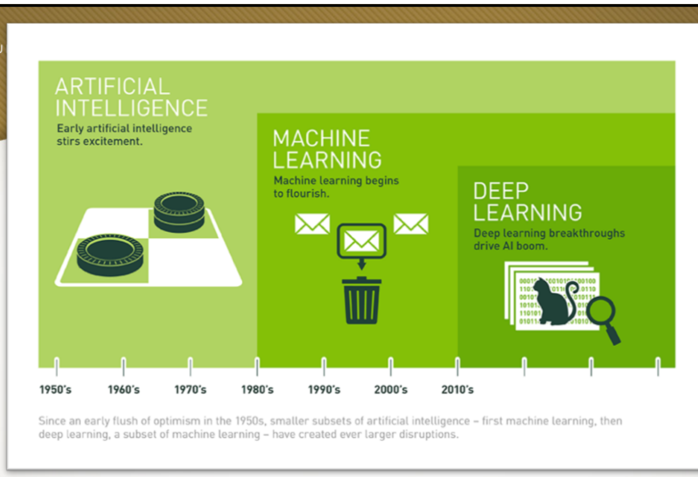

AI: The theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages

Machine Learning: Algorithms and statistical models that computer systems use to progressively improve their performance on a specific task.

Deep Learning: Learning data representations, as opposed to task-specific algorithms.

https://www.datasciencecentral.com/profiles/blogs/artificial-intelligence-vs-machine-learning-vs-deep-learning, accessed 12/10/2018

VANDERBILT UNIVERSITY School of Medicine

Measuring Classification: AUC ROC

- True Positive Rate = Sensitivity
- False Positive Rate = 1-Specificity
- AUC = Area under the curve
- ROC = Receiver operating characteristics

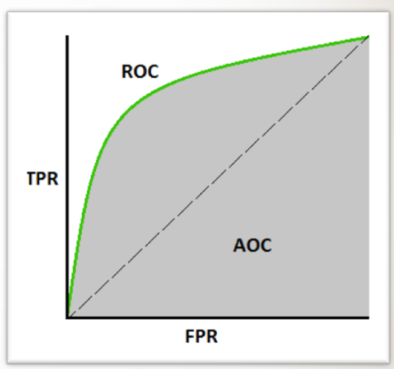

AUC=0.5 is chance separation
AUC=1 is perfect separation
AUC=0 has no separation

Roadmap

- Understand the context for the development of artificial intelligence in health care
- Explain how neural networks work
- Describe applications of AI in health care
  - Top 5 Technologies in Sleep and Anesthesia
- Describe limitations of the artificial intelligence

VANDERBILT UNIVERSITY School of Medicine

Early AI in Medicine

- Initial approach
  - Manipulation of symbolic expressions with rules of interference
  - Expert systems to automate reasoning process of experts

RULE 85

IF
1) The site of the culture is blood, and
2) The Gram stain of the organism is negative, and
3) The morphology of the organism is rod, and
4) The patient is a compromised host

THEN
There is suggestive evidence (0.6) that the identity of the organism is pseudomonas aeruginosa.

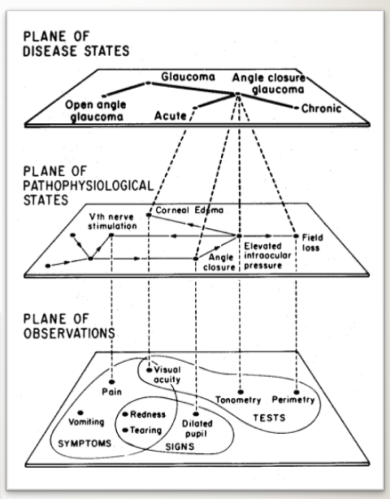

Shortliffe EH, Davis R, Axline SG, Buchanan BG, Green CC, Cohen SN. Comput Biomed Res. 1975 Aug;8(4):303-20.

VANDERBILT UNIVERSITY School of Medicine

How Actual Intelligence Works

- Layers of neuros
  - Information from outside world
  - Information transfer between neurons
    - Synaptic input can be strong or weak, connections can develop or be pruned
  - Transfer of information between layers of neurons

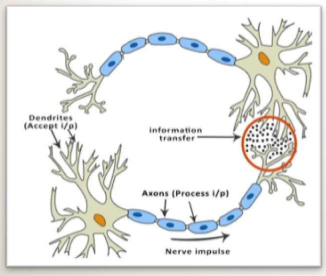

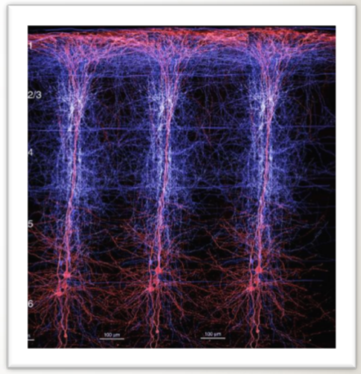

VANDERBILT UNIVERSITY School of Medicine

Neural Networks

- Artificial neural networks
  - Mimic neurons
  - Neural networks (1960s)
  - Back-propagation (1986)
    - Then computers got fast
- Hidden layers
  - Can represent dataset features
  - Number of hidden factors is specified
    - Inputs to those are not limited
- Convolutional neural network
  - Copies of features are made and averaged
  - Generalizes feature detection from one image area to another

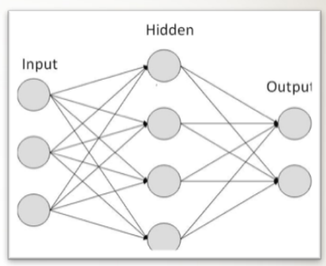

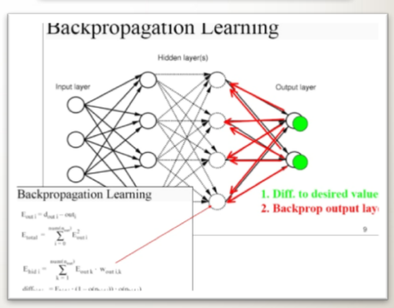

VANDERBILT UNIVERSITY School of Medicine

Logistic Regression and NN

$$\mathrm{logit}(p_i) = \ln\left(\frac{p_i}{1-p_i}\right) = \beta_0 + \beta_1 x_{1,i} + \cdots + \beta_M x_{m,}$$

Logistic regression is essentially a single-layer feed forward neural network

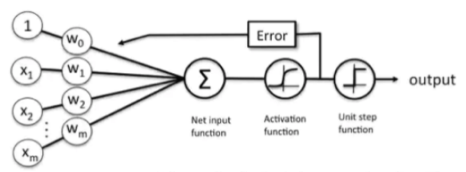

Schematic of a logistic regression classifier.

VANDERBILT UNIVERSITY School of Medicine

How to recognize a dog

- Training
  - Images with labels are fed to an untrained neural network
    - Wolves
    - Dogs
  - Weights are tweaked with back-propagation
- Validation
  - Images without labels are fed to network
    - Initial layer finds small features
    - Higher layer finds structural features
    - Highest lay represents concepts
  - Model performance is analyzed

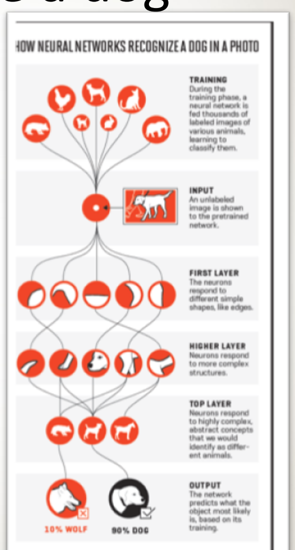

http://fortune.com/ai-artificial-intelligence-deep-machine-learning/, accessed 12/10/2018

VANDERBILT UNIVERSITY School of Medicine

Roadmap

- Understand the context for the development of artificial intelligence in health care
- Explain how neural networks work
- Describe applications of AI in health care
  - Top 5 Technologies in Sleep and Anesthesia
- Describe limitations of the artificial intelligence

VANDERBILT UNIVERSITY School of Medicine

Diabetic retinopathy

- Deep convolutional neural network
  - Trained on 128,175 retinal images
  - Separate validation datasets
- Images graded by 54 ophthalmologists
  - Diabetic retinopathy
  - Macular edema
  - Image quality
- Results
  - AUC 0.991, Sensitivity 93%, Specificity 97.5%

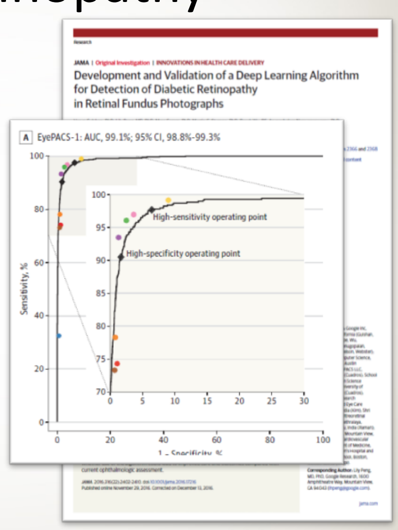

JAMA. 2016;316(22):2402-2410. doi:10.1001/jama.2016.17216

VANDERBILT UNIVERSITY School of Medicine

Breast Cancer Metastases

- Image data sets
  - 399 whole-slide images
    - 270 training
    - 129 test
  - Metastases: None, Macro, Micro
- Competition of
  - 32 machine learning algorithms
  - 11 pathologists
- Results
  - Varied from AUC 0.556 to 0.994
  - Top 5 algorithms were as good as pathologist without time constraints

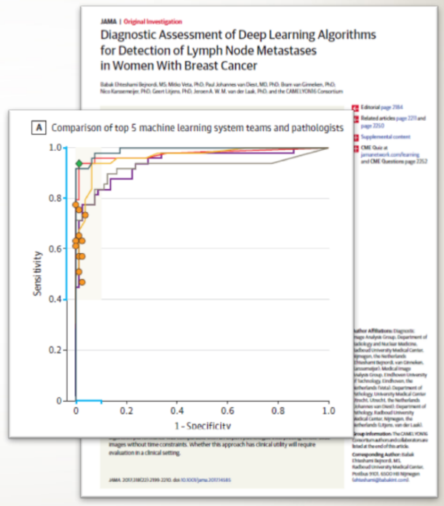

JAMA. 2017;318(22):2199-2210. doi:10.1001/jama.2017.14585

VANDERBILT UNIVERSITY School of Medicine

Melanoma Diagnosis

- Image data sets
  - 400 whole-slide images
    - 300 training
    - 100 test
  - Melanoma vs benign nevi
- Competition of
  - Google Inception v4 convoluted neural network
  - 58 pathologists
- Results
  - Neural network: AUC 0.86
  - Average pathologist: AUC 0.79

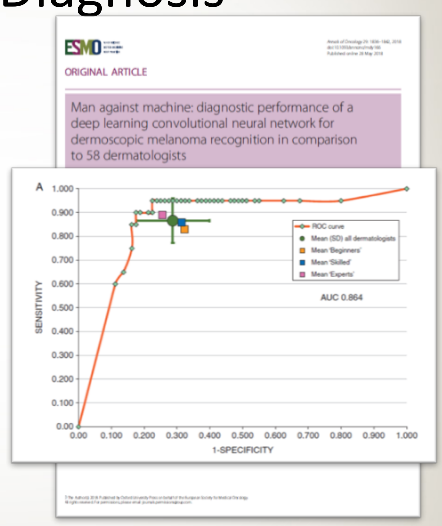

Ann Oncol. 2018 Aug 1;29(8):1836-1842. doi: 10.1093/annonc/mdy166.

VANDERBILT UNIVERSITY School of Medicine

Echo Interpretation

- Image data sets
  - 14,035 echo exams
- CNNs developed to
  - Identify viewpoints
  - Cardiac chamber segmentation in 5 views
  - Calculate volume, LV mass, EF
- Results for diagnostic accuracy
  - HOCM: AUC 0.93
  - Cardiac amyloidosis: 0.87
  - Pulmonary arterial HTN: 0.85

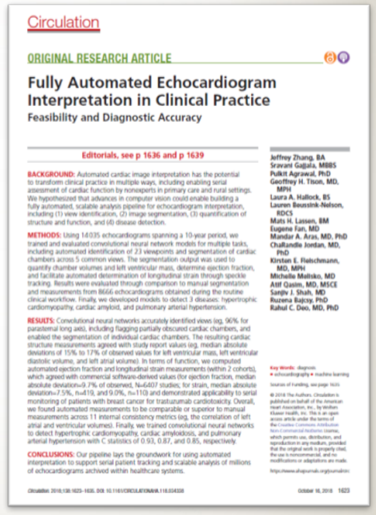
Circulation

ORIGINAL RESEARCH ARTICLE

Fully Automated Echocardiogram Interpretation in Clinical Practice

Feasibility and Diagnostic Accuracy

Editorials, see p 1636 and p 1639

Circulation. 2018;138:1623–1635. DOI: 10.1161/CIRCULATIONAHA

VANDERBILT UNIVERSITY School of Medicine

Stroke Detection

- Surely we're not using this machine learning stuff yet, right?

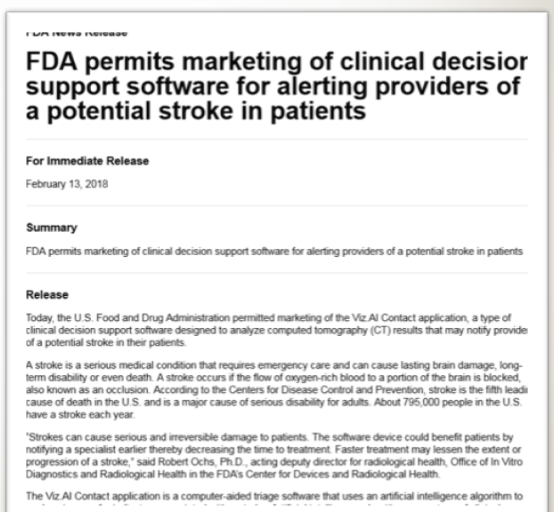
FDA permits marketing of clinical decision support software for alerting providers of a potential stroke in patients

For Immediate Release

February 13, 2018

Summary

FDA permits marketing of clinical decision support software for alerting providers of a potential stroke in patients

Release

Today, the U.S. Food and Drug Administration permitted marketing of the Viz.AI Contact application, a type of clinical decision support software designed to analyze computed tomography (CT) results that may notify provide of a potential stroke in their patients.

A stroke is a serious medical condition that requires emergency care and can cause lasting brain damage, long-term disability or even death. A stroke occurs if the flow of oxygen-rich blood to a portion of the brain is blocked, also known as an occlusion. According to the Centers for Disease Control and Prevention, stroke is the fifth leadi cause of death in the U.S. and is a major cause of serious disability for adults. About 795,000 people in the U.S. have a stroke each year.

"Strokes can cause serious and irreversible damage to patients. The software device could benefit patients by notifying a specialist earlier thereby decreasing the time to treatment. Faster treatment may lessen the extent or progression of a stroke," said Robert Ochs, Ph.D., acting deputy director for radiological health, Office of In Vitro Diagnostics and Radiological Health in the FDA's Center for Devices and Radiological Health.

The Viz.AI Contact application is a computer-aided triage software that uses an artificial intelligence algorithm to

VANDERBILT UNIVERSITY School of Medicine

Fracture Detection

- But probably just for that stroke stuff, no?

FDA News Release

FDA permits marketing of artificial intelligence algorithm for aiding providers in detecting wrist fractures

For Immediate Release

May 24, 2018

Summary

FDA permits marketing of artificial intelligence algorithm for aiding providers in detecting wrist fractures

Release

Today, the U.S. Food and Drug Administration permitted marketing of Imagen OsteoDetect, a type of computer-aided detection and diagnosis software designed to detect wrist fractures in adult patients.

"Artificial intelligence algorithms have tremendous potential to help health care providers diagnose and treat medical conditions," said Robert Ochs, Ph.D., acting deputy director for radiological health, Office of In Vitro Diagnostics and Radiological Health in the FDA's Center for Devices and Radiological Health. "This software can help providers detect wrist fractures more quickly and aid in the diagnosis of fractures."

The OsteoDetect software is a computer-aided detection and diagnostic software that uses an artificial intelligence algorithm to analyze two-dimensional X-ray images for signs of distal radius fracture, a common type of wrist fracture. The software marks the location of the fracture on the image to aid the provider in detection and diagnosis.

OsteoDetect analyzes wrist radiographs using machine learning techniques to identify and highlight regions of distal radius fracture during the review of posterior-anterior (front and back) and medial-lateral (sides) X-ray images of adult wrists. OsteoDetect is intended to be used by clinicians in various settings, including primary care, emergency medicine, urgent care and specialty care, such as orthopedics. It is an adjunct tool and is not intended to replace a clinician's review of the radiograph or his or her clinical judgment.

VANDERBILT UNIVERSITY School of Medicine

Automated Sleep Scoring

- Data sets
  - 10,000 clinical PSGs
  - 5804 research PSGs
- RCNNs developed to identify
  - Sleep stages
  - Sleep disordered breathing
  - Limb movements
- Results for diagnostic accuracy
  - 87.6%, 88.2%, 84.7%
  - All comparable to human experts
  - Almost as good with limited channels (i.e., at-home monitoring)

Journal of the American Medical Informatics Association, 25(12), 2018, 1643–1650
doi: 10.1093/jamia/ocy131
Advance Access Publication Date: 16 November 2018
Research and Applications

AMIA OXFORD

Research and Applications

Expert-level sleep scoring with deep neural networks

Siddharth Biswal,[1] Haoqi Sun,[2] Balaji Goparaju,[2,3] M Brandon Westover,[2,*] Jimeng Sun,[1,*] and Matt T Bianchi[2,3,*]

[1]School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA, [2]Neurology Department, Massachusetts General Hospital, Wang 720, Boston, MA, USA and [3]Division of Sleep Medicine, Harvard Medical School, Boston, MA, USA

*These authors contributed equally to this work

Corresponding Author: M. Brandon Westover, Massachusetts General Hospital, 55 Fruit Street, Boston, MA 02114, USA; mwestover@mgh.harvard.edu

Received 10 April 2018; Revised 17 August 2018; Editorial Decision 14 September 2018; Accepted 21 September 2018

ABSTRACT

Objectives: Scoring laboratory polysomnography (PSG) data remains a manual task of visually annotating 3 primary categories: sleep stages, sleep disordered breathing, and limb movements. Attempts to automate this process have been hampered by the complexity of PSG signals and physiological heterogeneity between patients. Deep neural networks, which have recently achieved expert-level performance for other complex medical tasks, are ideally suited to PSG scoring, given sufficient training data.
Methods: We used a combination of deep recurrent and convolutional neural networks (RCNN) for supervised learning of clinical labels designating sleep stages, sleep apnea events, and limb movements. The data for testing and training were derived from 10 000 clinical PSGs and 5804 research PSGs.
Results: When trained on the clinical dataset, the RCNN reproduces PSG diagnostic scoring for sleep staging, sleep apnea, and limb movements with accuracies of 87.6%, 88.2% and 84.7% on held-out test data, a level of performance comparable to human experts. The RCNN model performs equally well when tested on the independent research PSG database. Only small reductions in accuracy were noted when training on limited channels to mimic at-home monitoring devices: frontal leads only for sleep staging, and thoracic belt signals only for the apnea-hypopnea index.
Conclusions: By creating accurate deep learning models for sleep scoring, our work opens the path toward broader and more timely access to sleep diagnostics. Accurate scoring automation can improve the utility and efficiency of in-lab and at-home approaches to sleep diagnostics, potentially extending the reach of sleep expertise beyond specialty clinics.

VANDERBILT UNIVERSITY School of Medicine

Narcolepsy Diagnosis

- Data sets
  - 3,000 PSGs
- Multiple approaches to identify
  - Sleep stages
  - Type-1 Narcolepsy
- Results for accuracy
  - 87%
  - Better than any one scorer
  - For T1N, 96% specific, 91% sensitive

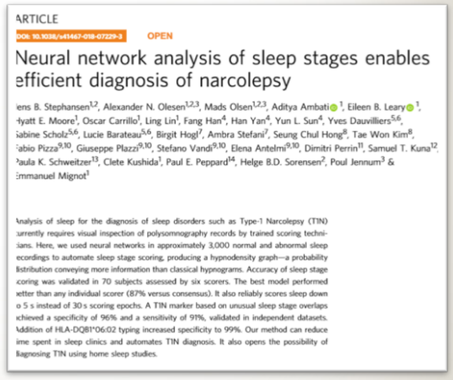

ARTICLE

DOI: 10.1038/s41467-018-07229-3 OPEN

Neural network analysis of sleep stages enables efficient diagnosis of narcolepsy

Jens B. Stephansen[1,2], Alexander N. Olesen[1,2,3], Mads Olsen[1,2,3], Aditya Ambati[1], Eileen B. Leary[1], Hyatt E. Moore[1], Oscar Carrillo[1], Ling Lin[1], Fang Han[4], Han Yan[4], Yun L. Sun[4], Yves Dauvilliers[5,6], Sabine Scholz[5,6], Lucie Barateau[5,6], Birgit Hogl[7], Ambra Stefani[7], Seung Chul Hong[8], Tae Won Kim[8], Fabio Pizza[9,10], Giuseppe Plazzi[9,10], Stefano Vandi[9,10], Elena Antelmi[9,10], Dimitri Perrin[11], Samuel T. Kuna[12], Paula K. Schweitzer[13], Clete Kushida[1], Paul E. Peppard[14], Helge B.D. Sorensen[2], Poul Jennum[3] & Emmanuel Mignot[1]

Analysis of sleep for the diagnosis of sleep disorders such as Type-1 Narcolepsy (T1N) currently requires visual inspection of polysomnography records by trained scoring technicians. Here, we used neural networks in approximately 3,000 normal and abnormal sleep recordings to automate sleep stage scoring, producing a hypnodensity graph—a probability distribution conveying more information than classical hypnograms. Accuracy of sleep stage scoring was validated in 70 subjects assessed by six scorers. The best model performed better than any individual scorer (87% versus consensus). It also reliably scores sleep down to 5 s instead of 30 s scoring epochs. A T1N marker based on unusual sleep stage overlaps achieved a specificity of 96% and a sensitivity of 91%, validated in independent datasets. Addition of HLA-DQB1*06:02 typing increased specificity to 99%. Our method can reduce time spent in sleep clinics and automates T1N diagnosis. It also opens the possibility of diagnosing T1N using home sleep studies.

VANDERBILT UNIVERSITY School of Medicine

Sleep Sound Signatures

- Smartphone sonar
  - Emits frequency-modulated sounds
  - Listens to reflections, tracks chest/abd movement
  - Identifies obstructive, central and hypopnea events
- Accuracy of ApneaApp
  - 37 patients, 296 hours of observation
  - 99.57%, 98.70, 95.33%

(correlation coefficient vs PSG)

Contactless Sleep Apnea Detection on Smartphones

Rajalakshmi Nandakumar, Computer Science and Engineering, University of Washington, rajaln@uw.edu

Shyamnath Gollakota, Computer Science and Engineering, University of Washington, gshyam@uw.edu

Nathaniel Watson M.D., UW Medicine Sleep Center, University of Washington, nwatson@uw.edu

Abstract – We present a contactless solution for detecting sleep apnea events on smartphones. To achieve this, we introduce a novel system that monitors the minute chest and abdomen movements caused by breathing on smartphones. Our system works with the phone away from the subject and can simultaneously identify and track the fine-grained breathing movements from multiple subjects. We do this by transforming the phone into an active sonar system that emits frequency-modulated sound signals and listens to their reflections; our design monitors the minute changes to these reflections to extract the chest movements.

VANDERBILT UNIVERSITY School of Medicine

Attack of the Sleep Apps!

- Sleep.ai
  - Track your snoring
  - Track your grinding

- One scientific app for Grinding, Snoring and Obstructive Sleep Apnea
- Proof that you grind or snore
- Send the sound samples to your dentist or doctor for official sleep bruxism diagnose

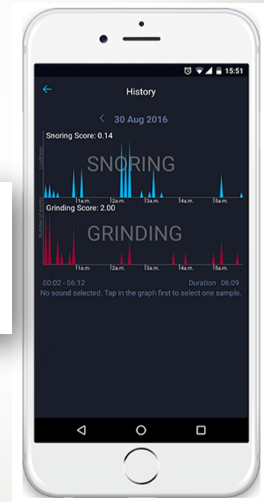

VANDERBILT UNIVERSITY School of Medicine

CPAP Compliance

- Somnoware
  - Prediction modeling for short and long-term compliance
  - Data fed from EHR
    - Demographics
    - Comorbidities
    - Prior labs
    - Questionnaire data

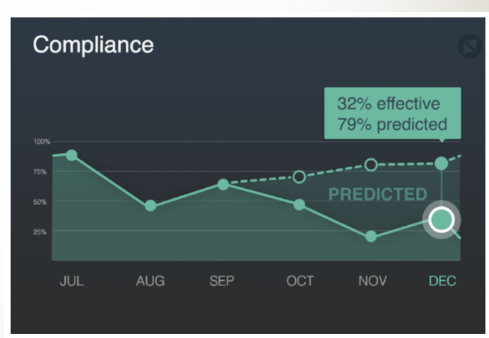

AI-Powered Mattress?!?

- Switching over to full AI hype mode…

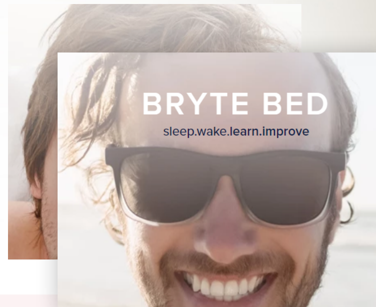

Are you sleeping on a Dead Bed?

The majority of beds used today are built on 19th century technology (static foam, springs, etc). Imagine the impact better sleep technology could have on your health, productivity and happiness.

We took an every day object and created an intelligent sleep platform that will transform your sleep experience and your life. (You'll never want to sleep on a dead bed again.)

VANDERBILT UNIVERSITY School of Medicine

AI-Powered Mattress?!?

BED FEATURE	DEAD BED	DEAD BED + BASIC TECH	BRYTE BED BRYTE FEATURES
Support Type	Foam, Springs or Air Bed	Foam, Springs or Air Bed	Hi-Res Dynamic Support
Dynamic Coils™ 100 Dynamic Coils into 16 zones with real-time adjustments			
Real Time Support Adjustments	-	Some	✔
Dynamic Support +			
Dynamic Dual Sides™ +			
AI-Based Sleep Management	-	-	✔
Personal Sleep Expert +			
Improves Over Time	-	-	✔
Personalized Learning +			
Relaxation Support	-	-	✔
Dynamic Lull™ +			
Nightly Bed & Room Preparation	-	-	✔

As low as $166/mo.
Learn more

QUEEN – $5,950

ADD TO CART

VANDERBILT UNIVERSITY School of Medicine

AI in Medicine

- Faster, cheaper computers have enabled cool technology
- Good algorithms exist for image interpretation and complex pattern recognition (i.e., PSG)
- Likely will make sleep studies more efficient
- Unclear if they will make your mattress work better

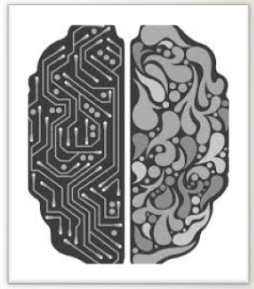

VANDERBILT UNIVERSITY School of Medicine

Roadmap

- Understand the context for the development of artificial intelligence in health care
- Explain how neural networks work
- Describe applications of AI in health care
  - Top 5 Technologies in Sleep and Anesthesia
- Describe limitations of the artificial intelligence

VANDERBILT UNIVERSITY School of Medicine

The Barrier of Meaning

- AI still doesn't understand the meaning of things
  - *"The bareheaded man needed a hat"*
  - *"The bear headed man needed a hat"*

https://www.nytimes.com/2018/11/05/opinion/artificial-intelligence-machine-learning.html, accessed 12/17/2018

VANDERBILT UNIVERSITY School of Medicine

Hackable

- Minor changes can disrupt AI-based systems

Using this methodology, we evaluate the efficacy of physical adversarial manipulations on real objects. With a perturbation in the form of only black and white stickers, we attack a real stop sign, causing targeted misclassification in 100% of the images obtained in lab settings, and in 84.8% of the captured video frames obtained on a moving vehicle(field test) for the target classifier

Robust Physical-World Attacks on Deep Learning Models. arXiv:1707.08945

VANDERBILT UNIVERSITY School of Medicine

Confusable

- AI-systems may not be able to deal with the unexpected

The Elephant in the Room. arXiv:1808.03305

VANDERBILT UNIVERSITY School of Medicine

The Black Box

- What hidden factors are being used?
  - Large numbers of complicated, weak connections
  - Not interpretable, in general
  - Initial weights are set randomly, so same NN with same data set will generate different features
  - No correct NN, just modeled relationship between inputs and outputs

AI Taking Over

- "People worry that computers will get too smart and take over the world, but the real problem is that they're too stupid and they've already taken over the world."
  - Pedro Domingos

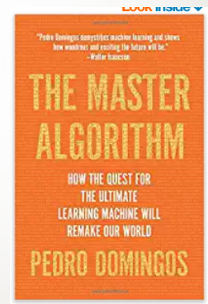

VANDERBILT UNIVERSITY School of Medicine

AI and Healthcare

- What's driving adoption of AI?
  - Digital imaging > Human interpretation
  - Digitization of health-related data and sharing
  - Adaptability of deep learning to heterogenous data
  - Capacity of deep learning for hypothesis generation
  - Potential to streamline workflow / empower patients
  - Rapid diffusion of deep learning tools
  - Better, faster technology

Stead WW. Clinical Implications and Challenges of Artificial Intelligence and Deep Learning. JAMA. 2018 Sep 18;320(11):1107-1108.

VANDERBILT UNIVERSITY School of Medicine

Are Robot Zombies Going To Take Over Medicine?

- Seems unlikely
  - Technology solves problems narrowly
  - No “generalized AI” on horizon
  - No capacity for reasoning
  - Output is really a statistical prediction
  - Doesn’t work well outside of trained data
- Physicians fundamentally trained to reason and act under uncertainty
- Best approach is to augment human intelligence in health care
  - AI = Augmented Intelligence

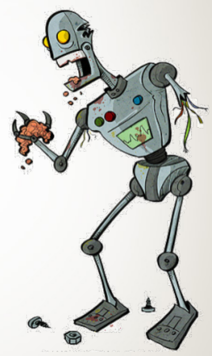

Stead WW. Clinical Implications and Challenges of Artificial Intelligence and Deep Learning. JAMA. 2018 Sep 18;320(11):1107-1108.

VANDERBILT UNIVERSITY School of Medicine

Thanks!

Questions?