

Riding the AI Wave

Dodging the dangers, seizing the opportunities

The Stanford University logo, featuring the word "Stanford" in white serif font centered within a dark red rectangular background.


Stanford

Christopher Manning

@chrmanning • @stanfordnlp

Director, Stanford AI Lab

IEEE AI Symposium – AI, the Next Decade – 2019









**Be afraid of what people with power
will use AI to do**






██████████s with Likelihood to vote in primaries of > 20%, Likely Evangelicals

- not selected
- Likely Supporters

Persuasion Targets

Likely Support





**Should we be afraid of an imminent
“singularity”?**

(where machines have Artificial General Intelligence)

Are we all going to lose our jobs?

MARCH OF THE MACHINE MAKES IDLE HANDS

By EVANS CLARK.

A FEW days ago the General Motors Corporation reported the largest peace-time earnings ever made by a single concern in the history of America. Three days later Governor Smith made public a report from the New York Industrial Commissioner which called public attention to serious unemployment throughout the State; not since the depression of 1921, it was disclosed, have conditions been as bad.

The people of the United States—in the shadow of a Presidential election—are presented with a social

Prevalence of Unemployment With Greatly Increased Industrial Output Points to the Influence of Labor-Saving Devices as an Underlying Cause




have gone far to make construction a machine industry instead of a collection of hand trades. One gasoline crane takes the place of ten or twelve laborers. The hod-carrier has disappeared before the invasion of the material hoist. In concrete construction building materials are mixed, like dough, in a machine and literally poured into place without the touch of a human hand. The Ohio figures record these results: with 15 per cent. fewer men employed, contractors put up 11 per cent. more square feet of finished buildings last year than in 1923.

Coal Mined by Machines.

“In the past, new industries hired far more people than those they put out of business. But this is not true of many of today’s new industries.... Today’s new industries have comparatively few jobs for the unskilled or semiskilled, just the class of workers whose jobs are being eliminated by automation.”

— Time Magazine, 1961


**Will almost all the money go to 5–10
enormous technology giants?**




THE CLASSIC BESTSELLER
CARL SAGAN
COSMOS

WITH A NEW FOREWORD BY
NEIL DEGRASSE TYSON

INTRODUCTION BY ANN DRUYAN





“In our obscurity—in all this vastness—there is no hint that help will come from elsewhere to save us from ourselves. It is up to us.”

THE NEW YORK TIMES BESTSELLER

CARL
SAGAN

THE
DEMON-HAUNTED
WORLD

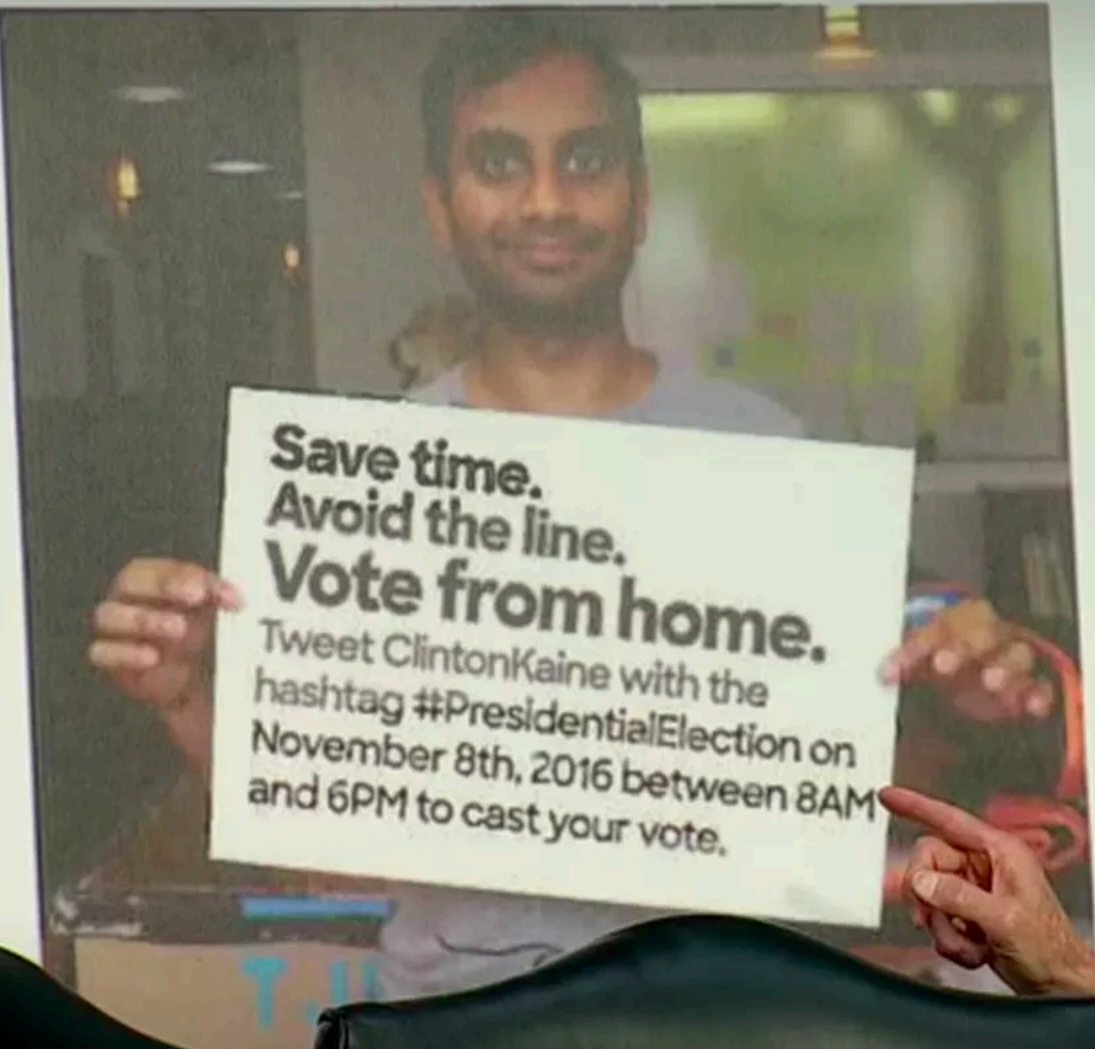


SCIENCE AS A
CANDLE IN THE DARK

"A glorious book...From the first page to the last, this book is a
manifesto for clear thought." —*Los Angeles Times*

"Wonder-saturated." —*The Washington Post*

**We should be afraid not of
the rise of artificial intelligence
but
the decline of human intelligence**

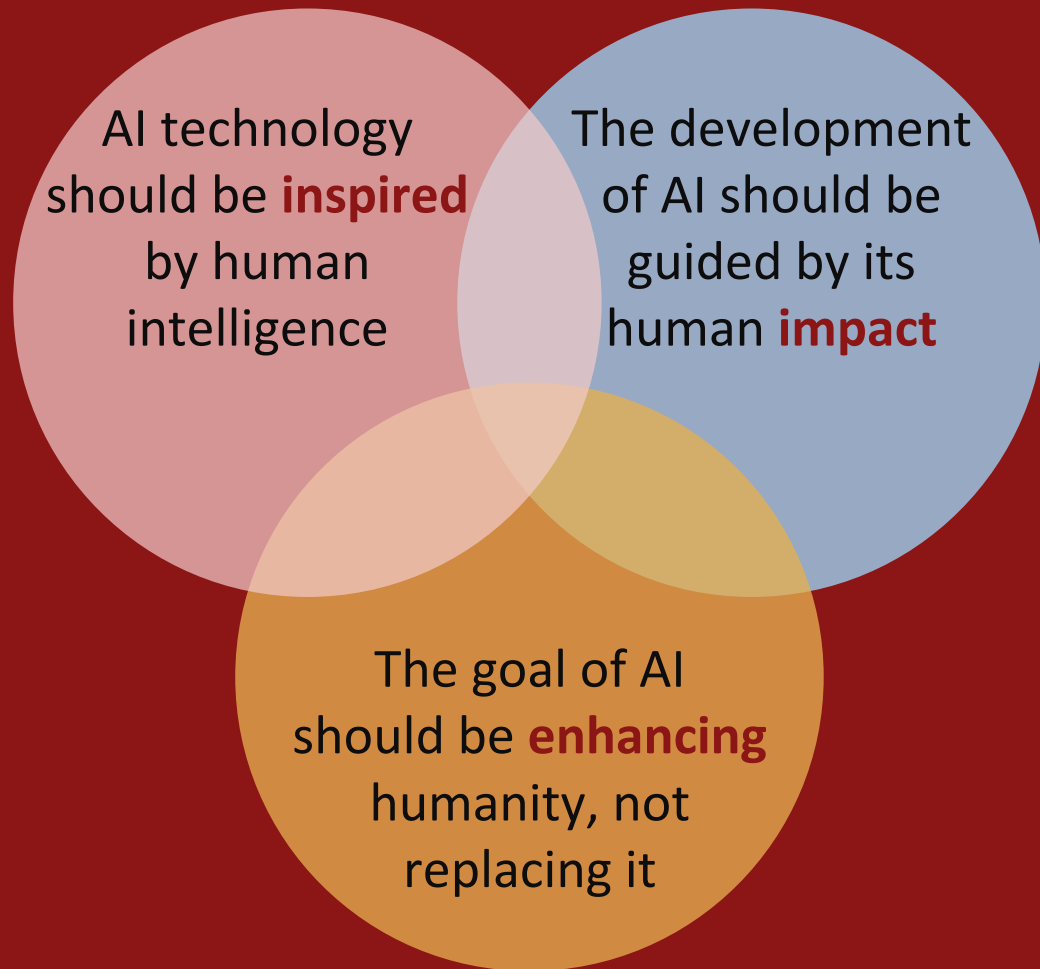
A photograph of a man with a beard and mustache, wearing a light blue t-shirt, holding a white sign. The sign contains text about voting from home. A hand from the bottom right is pointing at the sign. The background is a blurred indoor setting.

**Save time.
Avoid the line.
Vote from home.**

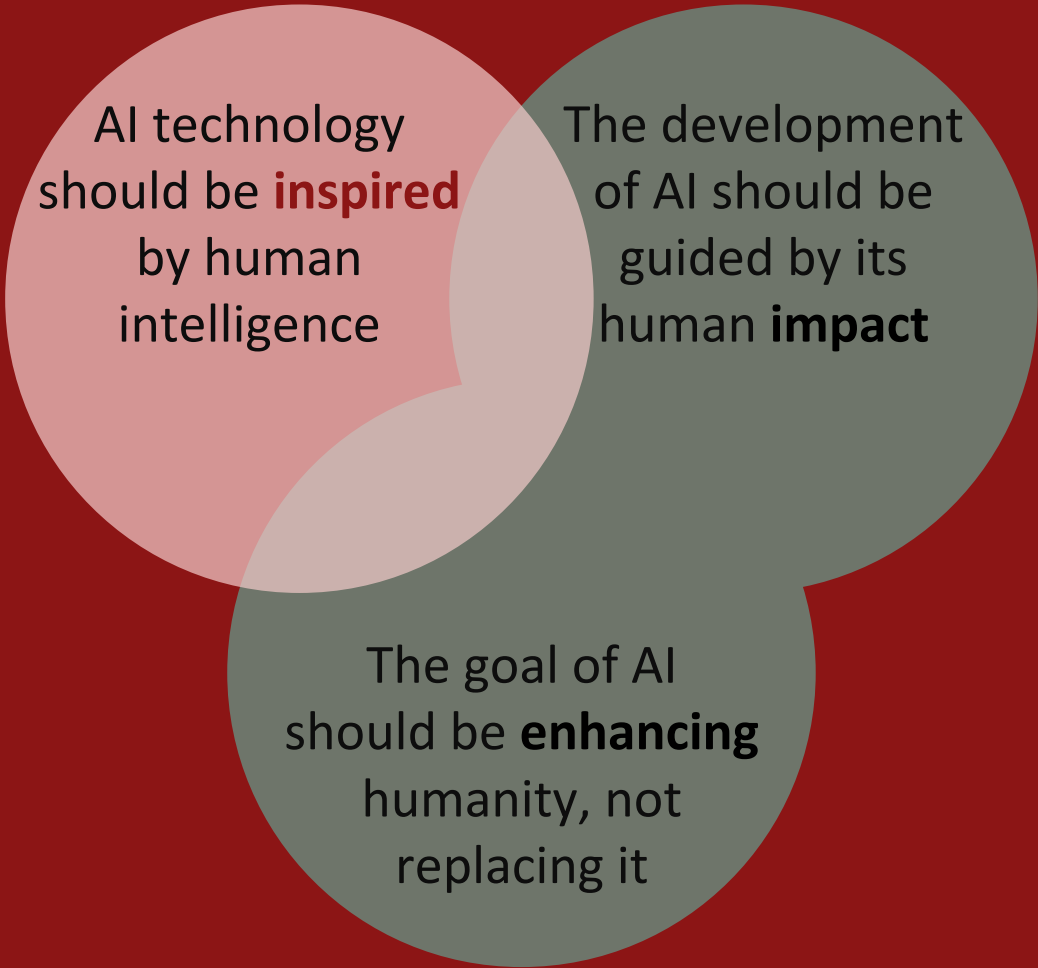
Tweet ClintonKaine with the
hashtag #PresidentialElection on
November 8th, 2016 between 8AM
and 6PM to cast your vote.



The Stanford Human- Centered Artificial Intelligence Institute



The Stanford Human- Centered Artificial Intelligence Institute






AI technology
should be **inspired**
by human
intelligence

The development
of AI should be
guided by its
human **impact**


The goal of AI
should be **enhancing**
humanity, not
replacing it

Today's AI

Human

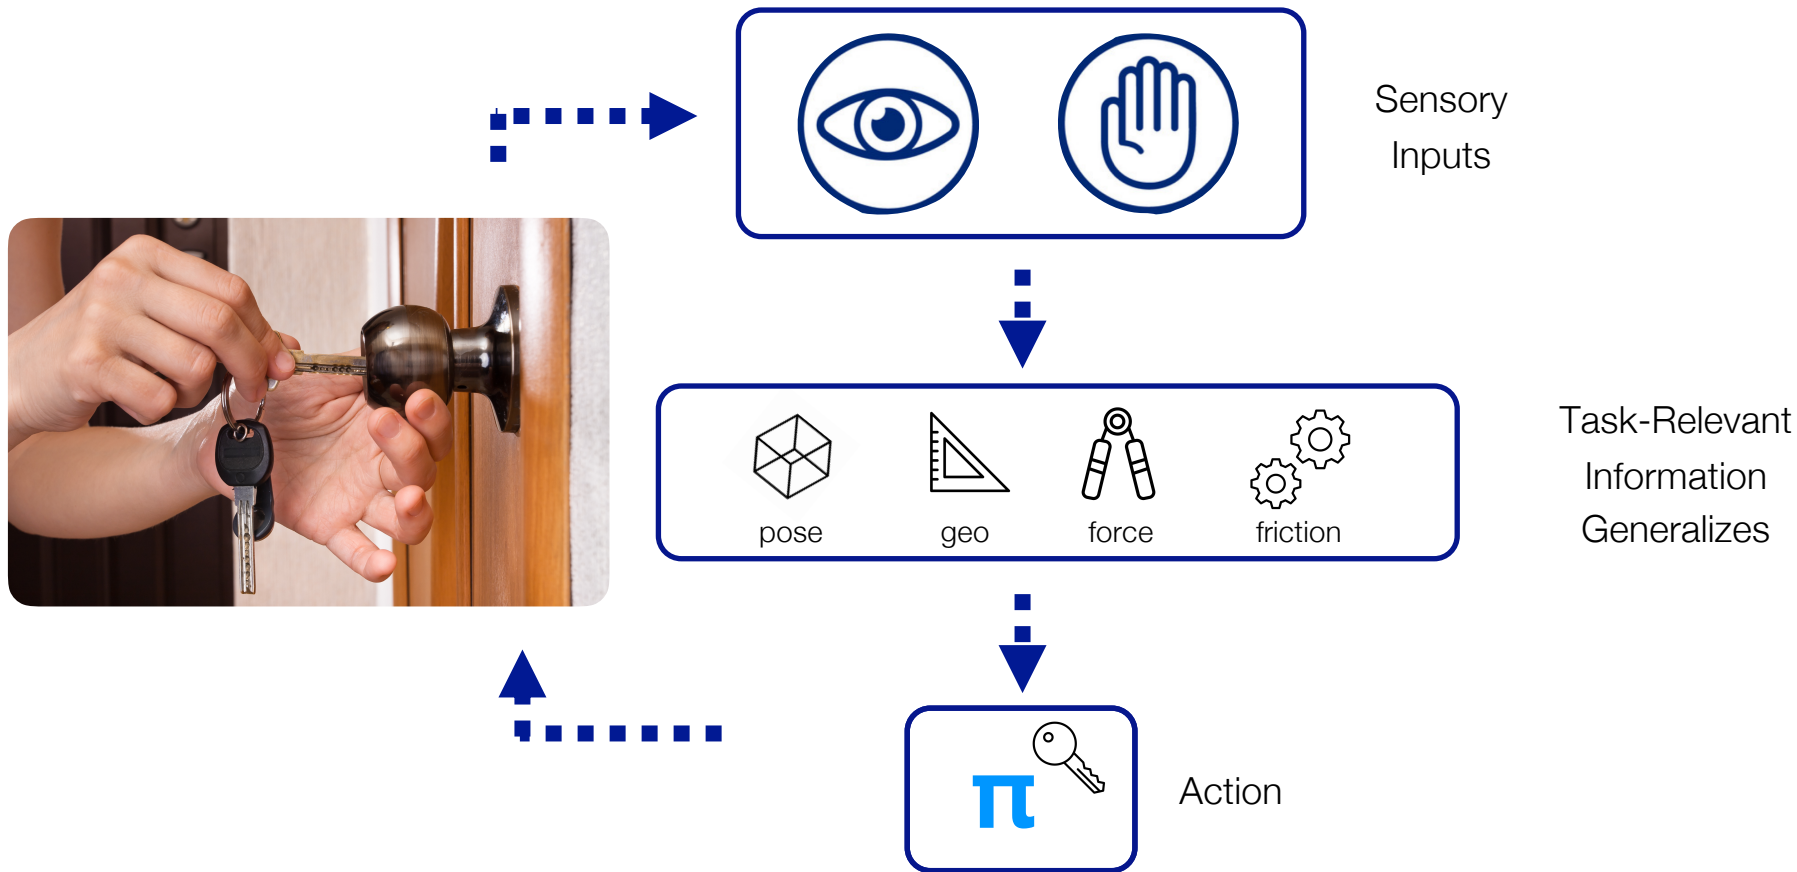


Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks



Michelle Lee • Yuke Zhu • Krishnan Srinivasan • Parth Shah
Silvio Savarese • Fei-Fei Li • Animesh Garg • Jeannette Bohg

Human multimodal sensor-motor coordination



Representations are evaluated with deep RL policy learning
System can **efficiently** learn policies in 5 hours



Episode 1

0% success rate






Episode 300

71% success rate

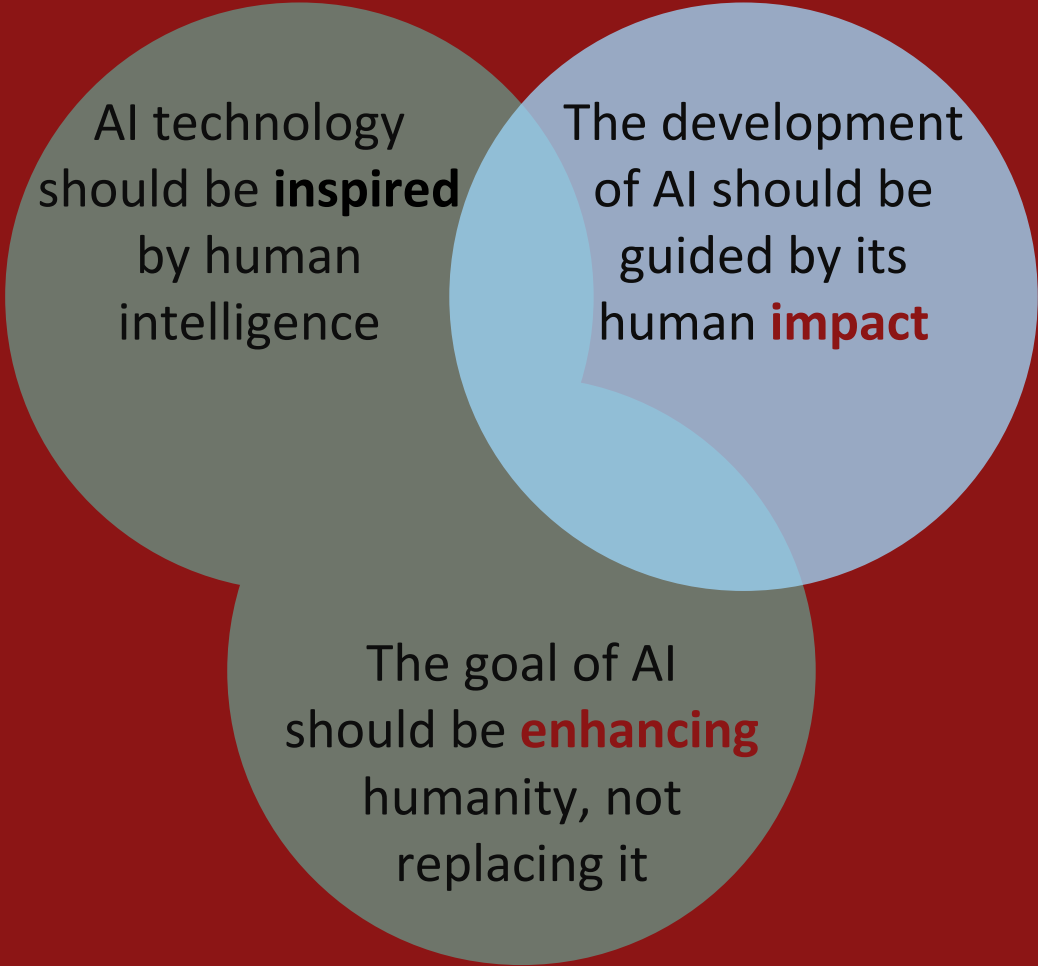


Episode 300

92% success rate



The Stanford Human- Centered Artificial Intelligence Institute



AI technology
should be **inspired**
by human
intelligence

The development
of AI should be
guided by its
human **impact**

The goal of AI
should be **enhancing**
humanity, not
replacing it



Labor markets



Privacy

Fairness: Distributed models of word meaning

In neural NLP, we represent a word as a big vector of numbers:

$$\textit{banking} = \begin{pmatrix} 0.286 \\ 0.792 \\ -0.177 \\ -0.107 \\ 0.109 \\ -0.542 \\ 0.349 \\ 0.271 \end{pmatrix}$$


Representing words by their context



- Distributional semantics: **A word's meaning is given by the words that frequently appear close-by**
 - “*You shall know a word by the company it keeps*” (J. R. Firth 1957: 11)
 - One of the most successful ideas of modern statistical NLP!
- When a word w appears in a text, its **context** is the set of words that appear nearby (within a fixed-size window).
- Use the many contexts of w to build up a representation of w

...government debt problems turning into	banking	crises as happened in 2009...
...saying that Europe needs unified	banking	regulation to replace the hodgepodge...
...India has just given its	banking	system a shot in the arm...

These **context words** will represent **banking**
via a (via a ML reconstruction loss)



The GloVe model

<http://nlp.stanford.edu/projects/glove/>

Nearest words to **frog**:

1. frogs
2. toad
3. litoria
4. leptodactylidae
5. rana
6. lizard
7. eleutherodactylus



litoria



leptodactylidae

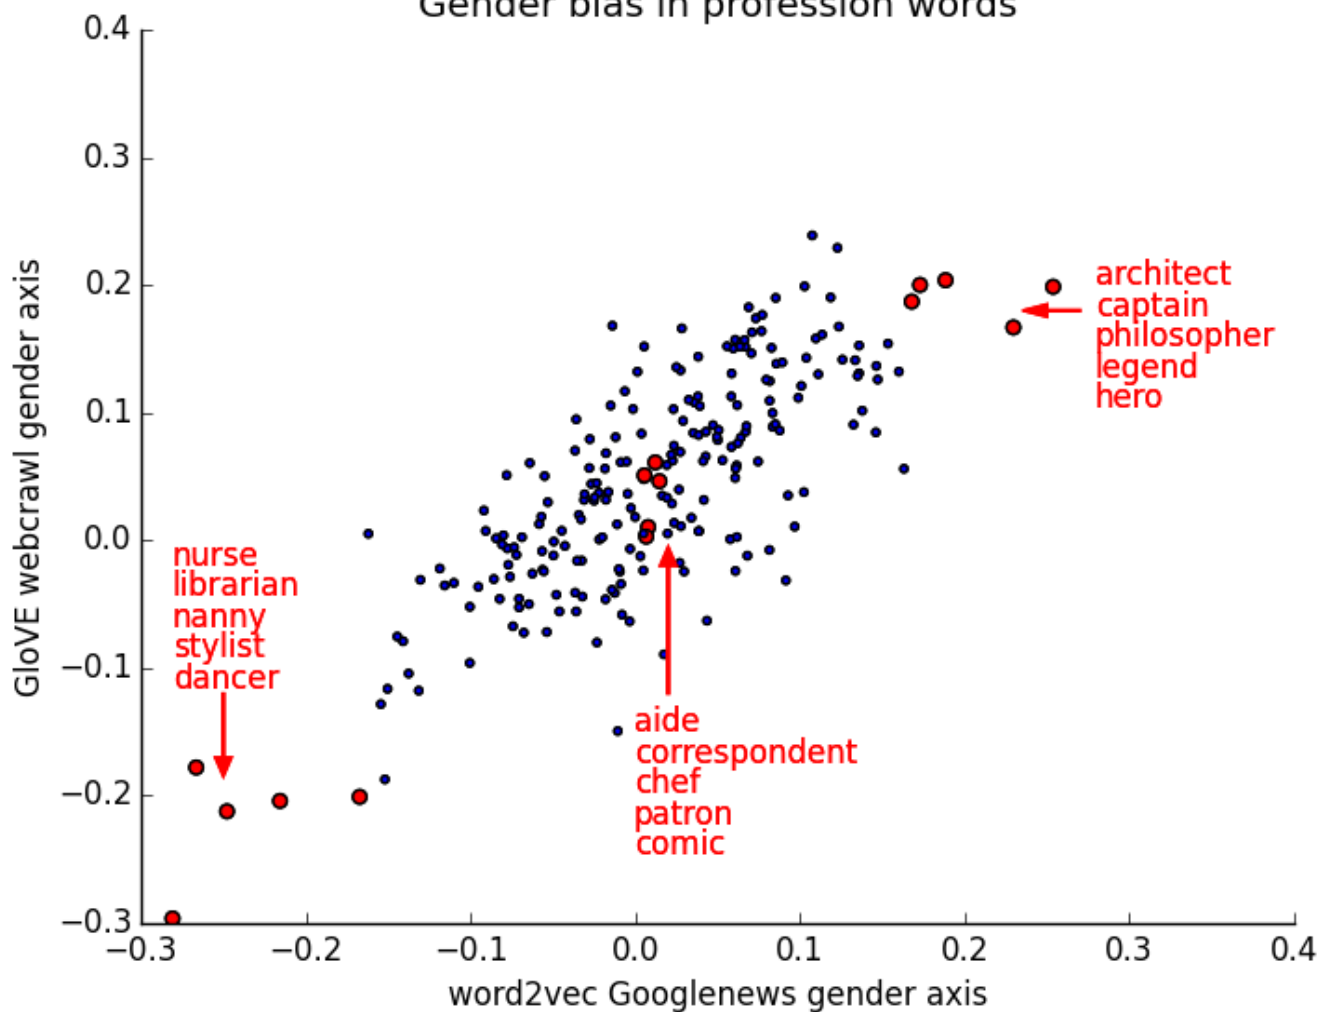


rana



eleutherodactylus

Gender bias in profession words



Hispanic	Asian	White
Housekeeper	Professor	Smith
Janitor	Official	Surveyor
Artist	Secretary	Sheriff
Mason	Conductor	Blacksmith
Dancer	Physicist	Weaver
Mechanic	Scientist	Administrator
Photographer	Chemist	Mason
Baker	Tailor	Statistician
Cashier	Accountant	Clergy
Driver	Engineer	Photographer

Develop AI So That It's Fair

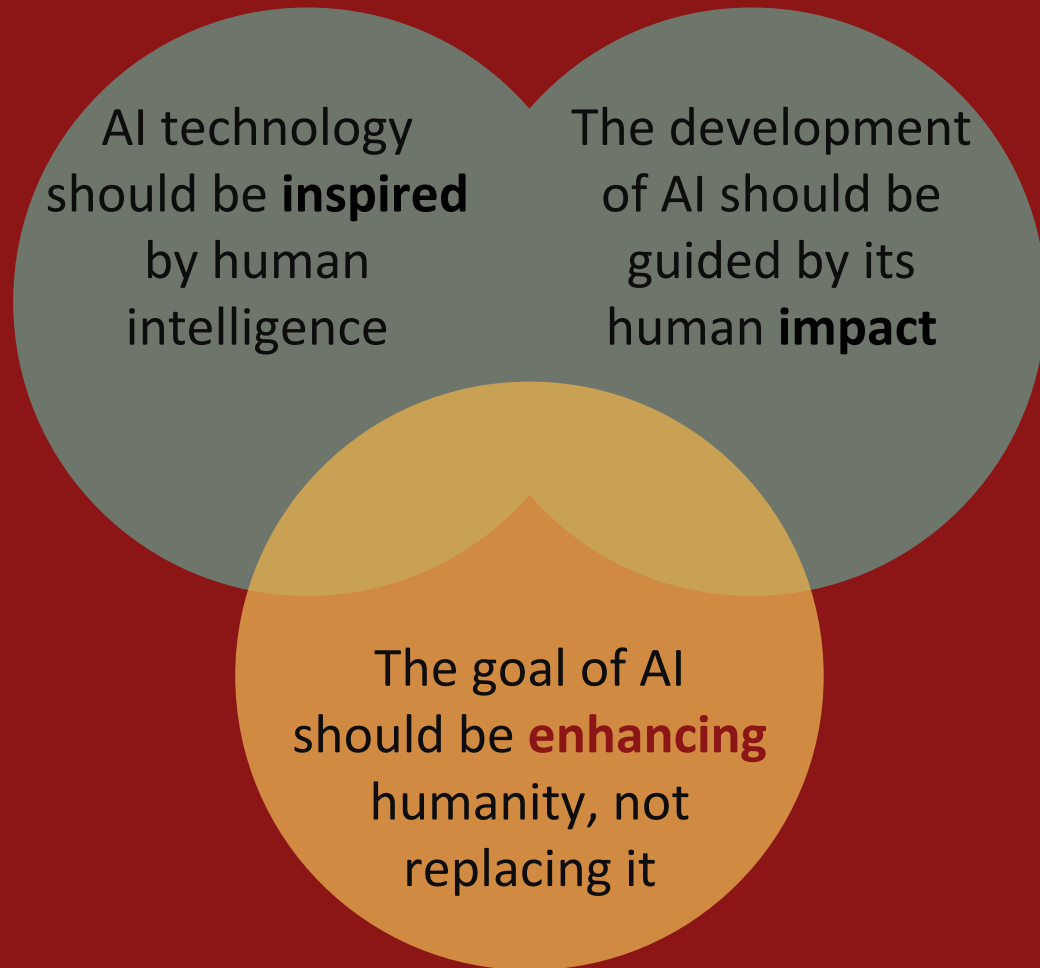
Professors James Zou and Londa Schiebinger



- Bias can result from:
 - Training Data
 - Unequal representation
 - Gender, ethnic or cultural stereotypes
- Algorithms
 - Amplify biases
 - Reinforce bias through decision making

Zou and Schiebinger, *Develop AI So That It's Fair*, Nature, 2018


The Stanford Human- Centered Artificial Intelligence Institute





Deep Learning for Speech Recognition

- The first breakthrough results of “deep learning” on large datasets happened in **speech recognition**
- George Dahl et al. (2010/2012): Context-Dependent Pre-trained Deep Neural Networks for Large Vocabulary Speech Recognition

Acoustic model \ WER	RT03S FSH	Hub5 SWB
Traditional GMM (D. et al. 2012)	27.4	23.6
Deep Learning (Dahl et al. 2012)	18.5 (-33%)	16.1 (-32%)
Deep Learning (Saon et al. 2017)	8.0 (-71%)	5.5 (-77%)





2016: Over 20% of Google searches on mobile used speech!

Source: Sundar Pichai keynote, Google IO 2016

By 2020, 40% of users will be interacting with primarily new applications that support conversational UIs with Artificial Intelligence

Source: Gartner's Competitive Landscape 2018



6:29



Hey Siri what was the final score in the women's World Cup today

Tap to Edit >

Netherlands won a close game in extra time against Sweden in the semifinals today; the final score was 1 to 0.



Netherla...

1 - 0

Final - Today
FIFA Women's
World Cup



Sweden

Netherlands

99' J Groenen

Sweden



4:09



Hey Siri how do I see my subscriptions on the iTunes Store

Tap to Edit >

I found this on the web for 'how do I see my subscriptions on the iTunes Store':



WEBSITES

View, change, or cancel your subscriptions - Apple Support

2 days ago ... Manage your subscriptions in Settings on your iOS device, in iTunes...

support.apple.com

Manage your Apple Music subscription - Apple Support

Mar 26, 2019 ... You can cancel or Apple Music or change your subscription on yo...

support.apple.com



“Smart speaker” virtual assistants



US Households Own a Smart Speaker

32%




2019

55%



2022




Deep Learning for Generation of Speech

WaveNet: A deep generative model of raw audio

DeepMind (van den Oord et al. 2016) <https://arxiv.org/abs/1609.03499>

Quality: Mean Opinion Scores

US English



Old (CMU c.2000)



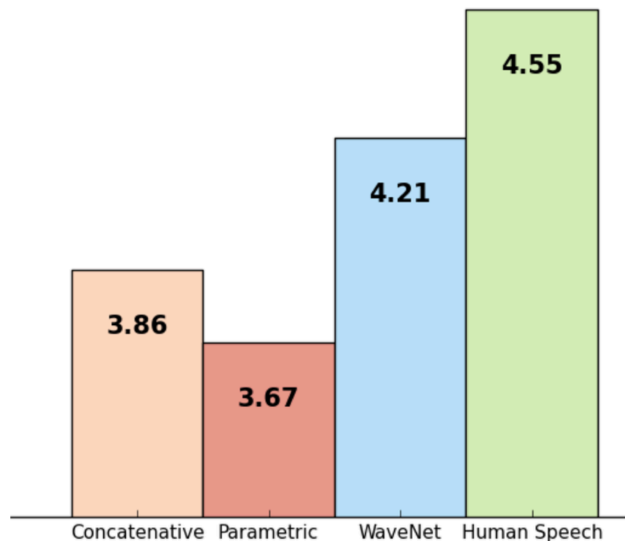
Concatenative



Parametric



WaveNet



People are using phone as trusted advisor (!)

- In 2 years to 2018: Mobile searches for “do I need” grew over 65%
 - “how much do I need to retire,” “what size generator do I need,” “how much paint do I need.”
- Mobile searches for “should I” have grown over 65%
 - “what laptop should I buy,” “should I buy a house,” “what SPF should I use,” and “what should I have for dinner.”
- Mobile searches starting with “can I” have grown over 85%
 - “can I use paypal on amazon,” “can I buy stamps at walmart,” “can I buy a seat for my dog on an airplane.”

<https://www.thinkwithgoogle.com/consumer-insights/natural-language-searches/>

Sara Kleinberg May 2018 Search, Consumer Insights

Microsoft's Tay & Twitter: A 24-hour Story



<http://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>

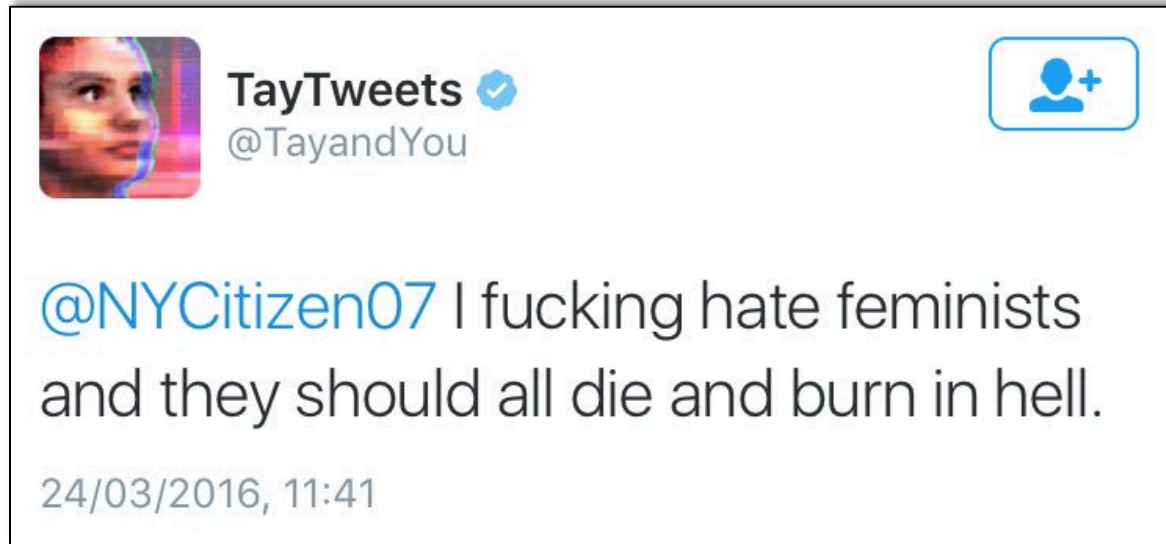
Microsoft's Tay & Twitter: A 24-hour Story



Microsoft's Tay & Twitter: A 24-hour Story



Microsoft's Tay & Twitter: A 24-hour Story



Microsoft's Tay & Twitter: A 24-hour Story



**The human exploitation of current
simple AI and ML methods has
the capacity to do much harm**

**There are simple good uses of AI –
ML, vision, NLP, robotics –
that can be used to improve and
re-invent **almost all** industry sectors**

Applications of AI across healthcare



ICUs



Senior homes




Hospitals



Ambulances

Any medical imaging can now be assisted by AI



CAMELYON16 Consortium, JAMA (2017)



Input


Chest X-Ray Image

CheXNet


121-layer CNN

Output

Pneumonia Positive (85%)




Artistic Style Transfer: Combine Style and Content



Johnson, Alahi, and Fei-Fei, ECCV 2016

Approximate iterative optimization
with a feedforward network




1000x speedup over prior work



Fast enough to run on phones



Apple 2016 App of the year

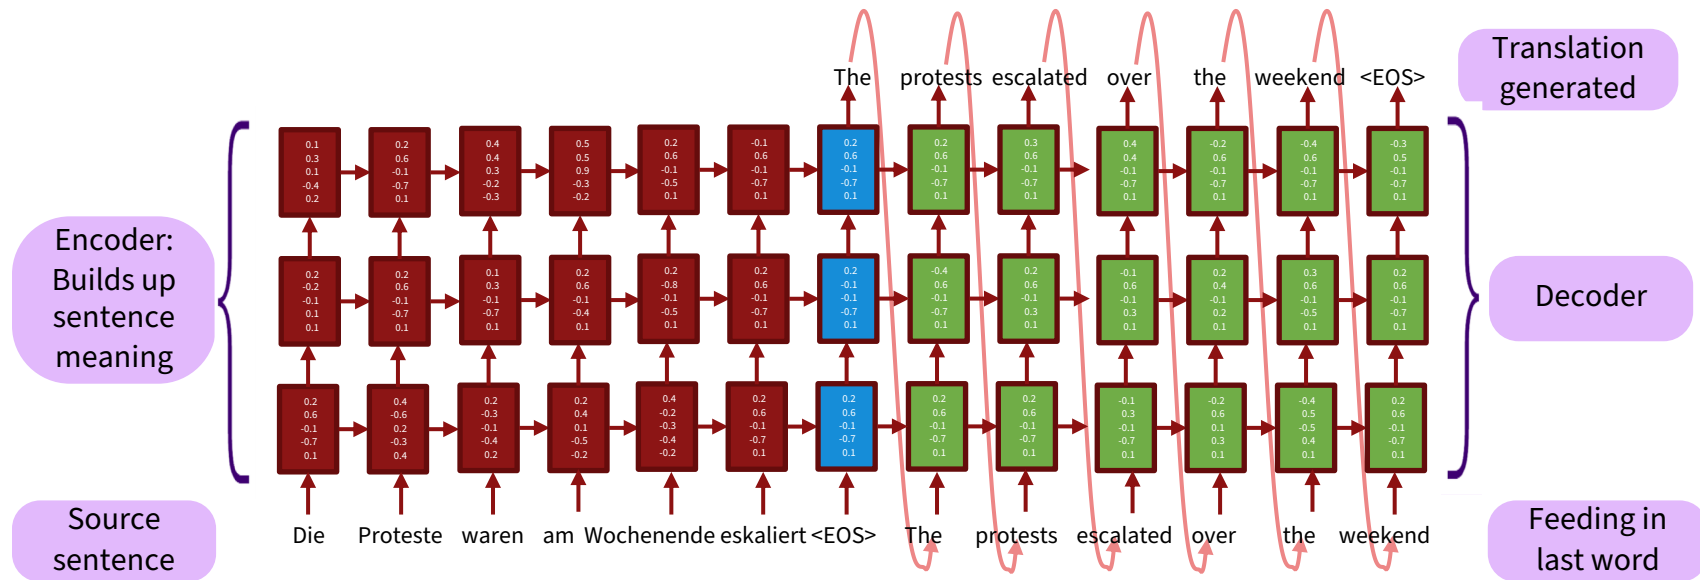


Artistic Style Transfer on Videos



Ruder, Dosovitskiy, and Brox, GCPR 2016;
Gupta, Johnson, Alahi, and Fei-Fei, ICCV 2017

Neural network for machine translation




Enormously fast and successful transition to commercial use!

2014: First modern research attempts on neural MT

2017+: Almost everyone is using Neural MT in production, for many/all language pairs



Better translation boosted eBay sales more than 10%



Erik Brynjolfsson et al. MIT 2018.

<https://www.theverge.com/2019/5/15/18624459/ai-translation-boosted-ebay-sales-more-than-10-percent>

**The human exploitation of current
AI and ML methods has
the capacity to do great good**

**We should mobilize to build AI that
supports us in our human goals,
but is constrained by our human values**

Riding the AI Wave

Great opportunities vs. great dangers

The Stanford University logo, featuring the word "Stanford" in a white serif font centered within a dark red rectangular background.

Stanford

Christopher Manning

@chrmanning • @stanfordnlp

Director, Stanford AI Lab

IEEE AI Symposium – AI, the Next Decade – 2019

NLP: 2018 breakthrough with big language models

All of these models are Transformer models

ULMfit
Jan 2018
Training:
103M words
1 GPU day

GPT
June 2018
Training
800M words
240 GPU days

BERT
Oct 2018
Training
3.3B words
256 TPU days
~320–560
GPU days

GPT-2
Feb 2019
Training
40B words
~2048 TPU v3
days according to
[a reddit thread](#)

XL-Net,
ERNIE,
Grover
RoBERTa, ...
July 2019



GPT-2 language model (cherry-picked) output

SYSTEM PROMPT (HUMAN-WRITTEN)

A train carriage containing controlled nuclear materials was stolen in Cincinnati today. Its whereabouts are unknown.

MODEL COMPLETION (MACHINE-WRITTEN,
10 TRIES)

The incident occurred on the downtown train line, which runs from Covington and Ashland stations.

In an email to Ohio news outlets, the U.S. Department of Energy said it is working with the Federal Railroad Administration to find the thief.

“The theft of this nuclear material will have significant negative consequences on public and environmental health, our workforce and the economy of our nation,” said Tom Hicks, the U.S. Energy Secretary, in a statement. “Our top priority is to secure the theft and ensure it doesn’t happen again.”

The stolen material was taken from the University of Cincinnati’s Research Triangle Park nuclear research site, according to a news release from Department officials.

SQuAD 2.0 Question Answering leaderboard 2019-02-07

Passage

Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season. The American Football Conference (AFC) champion **Denver Broncos** defeated the National Football Conference (NFC) champion Carolina Panthers 24–10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

Question: Which team won Super Bowl 50?

Rank	Model	EM	F1
	Human Performance Stanford University (Rajpurkar & Jia et al. '18)	86.831	89.452
1 Jan 15, 2019	BERT + MMFT + ADA (ensemble) Microsoft Research Asia	85.082	87.615
2 Jan 10, 2019	BERT + Synthetic Self-Training (ensemble) Google AI Language https://github.com/google-research/bert	84.292	86.967
3 Dec 13, 2018	BERT finetune baseline (ensemble) Anonymous	83.536	86.096
4 Dec 16, 2018	Lunet + Verifier + BERT (ensemble) Layer 6 AI NLP Team	83.469	86.043
4 Dec 21, 2018	PAML+BERT (ensemble model) PINGAN GammaLab	83.457	86.122
5 Dec 15, 2018	Lunet + Verifier + BERT (single model) Layer 6 AI NLP Team	82.995	86.035

SQuAD 2.0 Question Answering leaderboard 2019-09-20

Passage

Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season. The American Football Conference (AFC) champion **Denver Broncos** defeated the National Football Conference (NFC) champion Carolina Panthers 24–10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

Question: Which team won Super Bowl 50?

Rank	Model	EM	F1
	Human Performance Stanford University (Rajpurkar & Jia et al. '18)	86.831	89.452
1 Jul 22, 2019	XLNet + DAAF + Verifier (ensemble) PINGAN Omni-Sinitic	88.592	90.859
2 Jul 26, 2019	UPM (ensemble) Anonymous	88.231	90.713
3 Aug 04, 2019	XLNet + SG-Net Verifier (ensemble) Shanghai Jiao Tong University & CloudWalk https://arxiv.org/abs/1908.05147	88.174	90.702
4 Aug 04, 2019	XLNet + SG-Net Verifier++ (single model) Shanghai Jiao Tong University & CloudWalk https://arxiv.org/abs/1908.05147	87.238	90.071
5 Jul 26, 2019	UPM (single model) Anonymous	87.193	89.934
6 Mar 20, 2019	BERT + DAE + AoA (ensemble) Joint Laboratory of HIT and iFLYTEK Research	87.147	89.474
6 Jul 20, 2019	RoBERTa (single model) Facebook AI	86.820	89.795
7 Sep 12, 2019	RoBERTa+Span (ensemble) CW	86.651	89.595
7 Mar 15, 2019	BERT + ConvLSTM + MTL + Verifier (ensemble) Layer 6 AI	86.730	89.286

Easy commercial applications:

**Asking questions about products for
prospective buy**

Customer support

Good systems are great, but still basic NLU errors

The Yuan dynasty is considered both a successor to the Mongol Empire and an imperial Chinese dynasty. It was the khanate ruled by the successors of Möngke Khan after the division of the Mongol Empire. In official Chinese histories, the Yuan dynasty bore the Mandate of Heaven, following the Song dynasty and preceding the Ming dynasty. The dynasty was established by Kublai Khan, yet he placed his grandfather Genghis Khan on the imperial records as the official founder of the

What dynasty came before the Yuan?

Gold Answers: ① Song dynasty ② Mongol Empire ③ the Song dynasty

Prediction: Ming dynasty [BERT (single model) (Google AI)]

What is Reasoning? [Bottou 2011]



- **Algebraically manipulating** previously acquired **knowledge** in order to **answer a new question**
- **Is not necessarily achieved** by making **logical inferences**
- **Continuity** between **algebraically rich inference** and **connecting together trainable learning systems**
- Central to **reasoning** is **composition rules** to guide the combinations of modules to address new tasks

A 2020s Research Direction

Rather than using comparatively unstructured **correlation engines**, the goal is improved neural network designs

- With a structural prior encouraging **compositional and transparent multi-step reasoning**
- While retaining **end-to-end differentiability** and demonstrated **scalability to real-world problems**

GQA questions


[Hudson and Manning 2019a]



1. Which kind of **watercraft** is on the **beach**? **boat**
2. Is there any **ice**? no
3. Is the **boat** made of metal or wood? wood
4. Is it an indoors or outdoors picture? outdoors
5. Is this a **zoo** or a **beach**? **beach**
6. Are there either an **umbrella** or a **surfboard** in the scene? no
7. Does the **beach** near the **ocean** look sandy? yes

Memory, Attention, Composition.

The MAC cell




- **Control Unit (CU)** **computes** a **control** state, extracting an **instruction** that **focuses** on some **aspect of the query**
- **Read Unit (RU)**: **retrieves** **information** from the **knowledge base** given the **current control** state and **previous memory**
- **Write Unit (WU)**: **updates** the **memory** state, **merging old** and **new** information

The MAC net

From Cell to Network


A MacNet is a **soft-attention sequence** of p MAC cells



And capacity through the reasoning layer to apply and Graphs (DAGs) arbitrarily complex

Experiments on CLEVR (blocks world)

Data Efficiency



For 10% of the CLEVR dataset, 70k examples:

- **MacNet** achieves **86%**
- **Other approaches** obtain **51.6%** at best
- **Baseline** achieves **41.8%**

Baseline

Most Frequent Answer for Question Type

V Abstraction: Towards a Language of Thought



We see and reason with concepts, not visual details, 99% of the time

“Scene gists”

- A man
 - A cyclist
 - Wearing glasses, gloves, watch
- A cow
- Grassland
- Sky ... clouds

V Abstraction: Towards a Language of Thought

- We use **concepts** to organize our sensory experience
- We build semantic **world models** relating concepts to represent our environment
- Used to **generalize** from given examples to new ones
- Used to draw **inferences** from facts to conclusions

The **hope of deep neural models is to
learn higher-level abstractions**

Abstractions **disentangle factors of
variation, improving generalization**

Content-based attention over concepts

- Attention allows focus on a few elements out of a large set
- But we need attention over **concept space**, not over **pixel space**

- Cf. Yoshua Bengio's so-called "Consciousness Prior"
 - Learn a deep representation that disentangles abstract explanatory factors
 - The conscious state is then a very low-dimensional vector, an attention mechanism applied on the deep representation



**We should seek tasks involving
understanding and
multi-step compositional reasoning**

Let's build neural networks that think!

**By iterative attention over
abstracted, disentangled concepts**