



AI / ML research group Health Informatics

Jim Warren

*Professor of Health Informatics, School of Computer Science,
University of Auckland*



AI 'outperforms' doctors diagnosing breast cancer




Fergus Walsh
Medical correspondent
@BBCFergusWalsh

🕒 2 January 2020



🔗 Share



GETTY IMAGES

AI was as accurate as two doctors working together

Google's AlphaGo beats top human (May, 2017)




Why is the Go victory so impressive?



1997

- Branching factor
 - Number of options at each move
 - Rate of 'combinatorial explosion' limiting depth of brute force look-ahead
- Used deep learning methods
 - Multi-layer artificial neural networks
 - Big Data
 - Massive computation

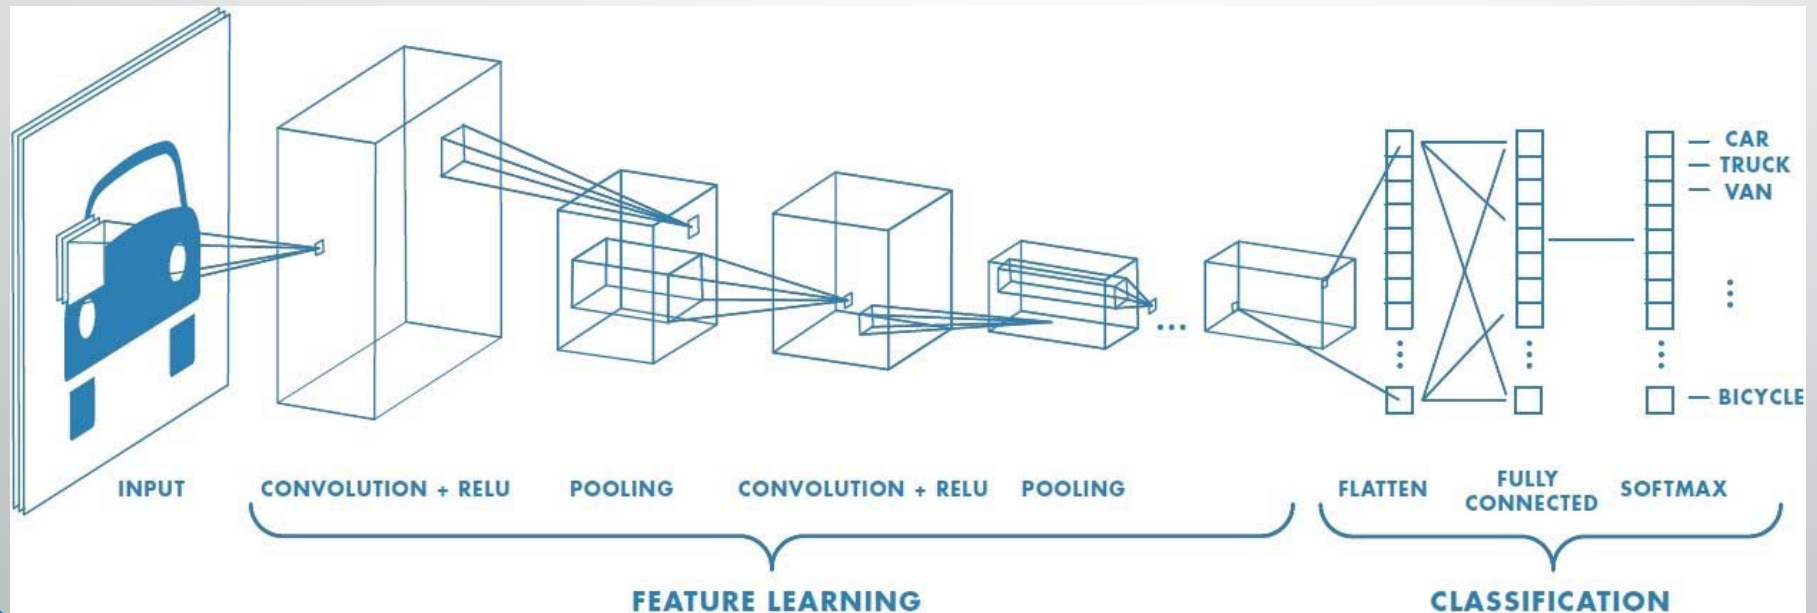
What is an artificial neural network



“Deep Learning”

Need:

- Lots of data
- Lots of compute power (often cloud based)



Convolutional neural network (from towardsdatascience.com)

Diabetic retinopathy classifier

- Detecting diabetic retinopathy as compared to panel of US licensed ophthalmologists and ophthalmology senior residents for 128,000 retinal images*
 - 97.5% sensitivity, 93.4% specificity
- Limitations
 - Algorithm wouldn't necessarily detect non-diabetic retinopathy lesions (outside its training data) nor would it be a replacement for a comprehensive eye exam
- 2352 citations (21 July 2020)

* Gulshan V, Peng L, Coram M, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA 2016;**316**(22):2402-10.



*aka, Digital Health,
L'informatique Médicale*

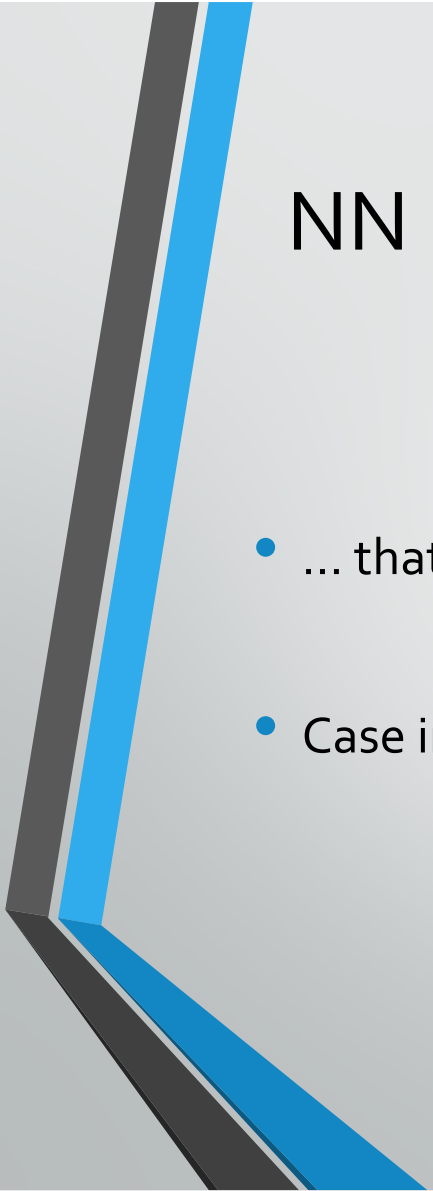
'Health Informatics'

- **Management and use of patient health care information**
- Multidisciplinary field aiming to improve healthcare delivery with IT in terms of quality, efficiency and/or creating new opportunities
- Includes AI in Medicine ('expert system', 'decision support system'), boring information systems stuff 😊 (implementing large-scale organizational change), and use of IT by health consumers
- Can be seen to also include Bioinformatics (computing about genes, proteins etc.)

Are doctors going the way of the buggy whip?


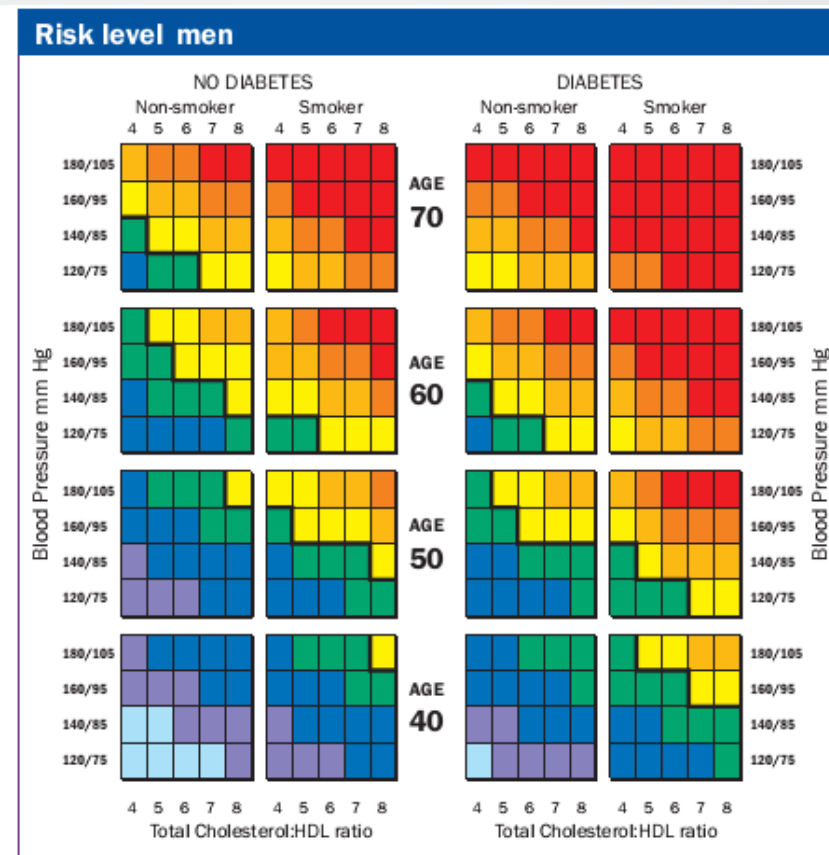
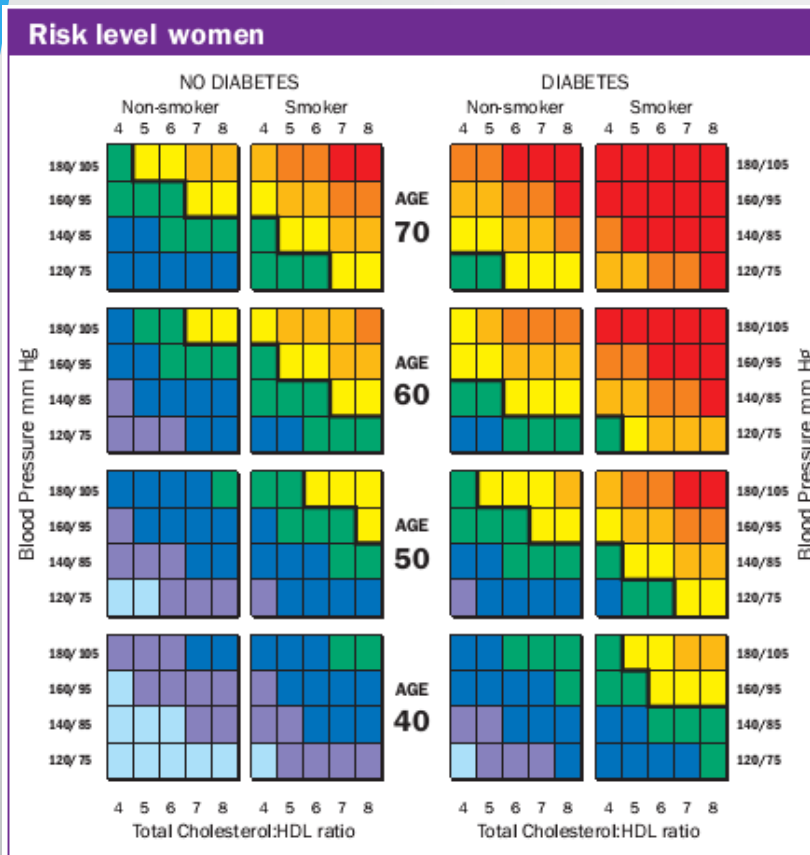
- We can replace them with machines at this point, right?
- What bright person would head to med school now?





NN based image classifiers aren't so different from some existing systems...

- ... that have integrated with the healthcare system just fine
- Case in point: the PREDICT CVD risk decision support system



How to use the Tables

- Identify the table relating to the person's sex, diabetic status, smoking history and age.
- Within the table choose the cell nearest to the person's age, blood pressure and TC:HDL ratio. When the systolic and diastolic values fall in different risk levels, the higher category applies.
- For example, the lower left cell contains all non-smokers without diabetes who are less than 45 years and have a TC:HDL ratio less than 4.5 and a blood pressure less than 130/80 mm Hg. People who fall exactly on a threshold between cells are placed in the cell indicating higher risk.

Figure 2: Assessing 5-year cardiovascular risk and treatment benefit

What's behind the prediction?

- Logistic regression
 - Log of the odds of an outcome (e.g. a cardiovascular disease event, such as a heart attack) as a weighted function of a number of risk factors (blood pressure, smoking, cholesterol, etc.)

$$\text{logit}(\mathbb{E}[Y_i \mid x_{1,i}, \dots, x_{m,i}]) = \text{logit}(p_i) = \ln \left(\frac{p_i}{1 - p_i} \right) = \beta_0 + \beta_1 x_{1,i} + \dots + \beta_m x_{m,i}$$

- Weights are learned by fitting to population health data
- For the scientific mind, seeing the 95% confidence interval of a Beta may be the way to go, but most people will appreciate the graphics

IT has been disrupting for a while in the form of 'Consumer Health Informatics'

- The Web (since 1993)
 - Anybody can come to a doctor's office brandishing the latest research (or just some junk that from a crackpot's blog, or something a drug company has put on the Web)
- PatientsLikeMe
 - Patients share experiences, and data, on off-label treatment for their condition (self-organized trials!)
- 23andme
 - Direct-to-patient genetic test results

Doctor's authority
challenged

Choice of research topics taken
out of hands of academia and
research councils

Doctor taken out of
the loop

Going directly to the computer for treatment: SPARX – gamified self-help

- Youth self-help for depression and anxiety as a first-person adventure videogame

Applies well-established methods of Cognitive Behaviour Therapy (CBT), making them more interesting with gamification

E.g. non-player characters and a companion bird called Hope; mini-games, exploration



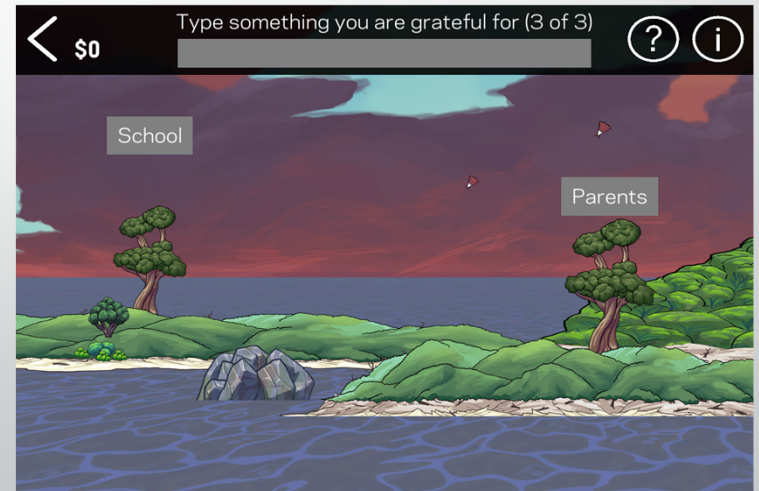
SPARX: evidence and implementation

- Not worse than usual care* (face-to-face counselling)
- Available for public use in New Zealand
 - <https://www.sparx.org.nz/>
- Now available reformatted as a mobile app
- SPARX exemplifies a scalable addition to the existing healthcare system
 - Available when a counsellor or other source of support may not be
 - Avoids some of the stigma of seeking help with mental health

* SN Merry, K Stasiak, M Shepherd, C Frampton, T Fleming, MFG Lucassen. The effectiveness of SPARX, a computerised self help intervention for adolescents seeking help for depression: randomised controlled non-inferiority trial. *Bmj* 344, e2598, 2012.



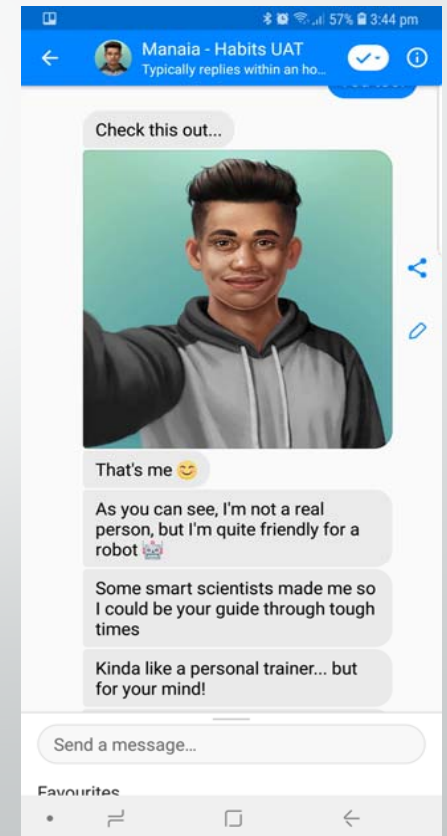
The Quest – *Te Whitianga*



Remind, Relax, Reengerise, Rethink, Resolve, Relate

HeadStrong – a dialog agent for NZ youth

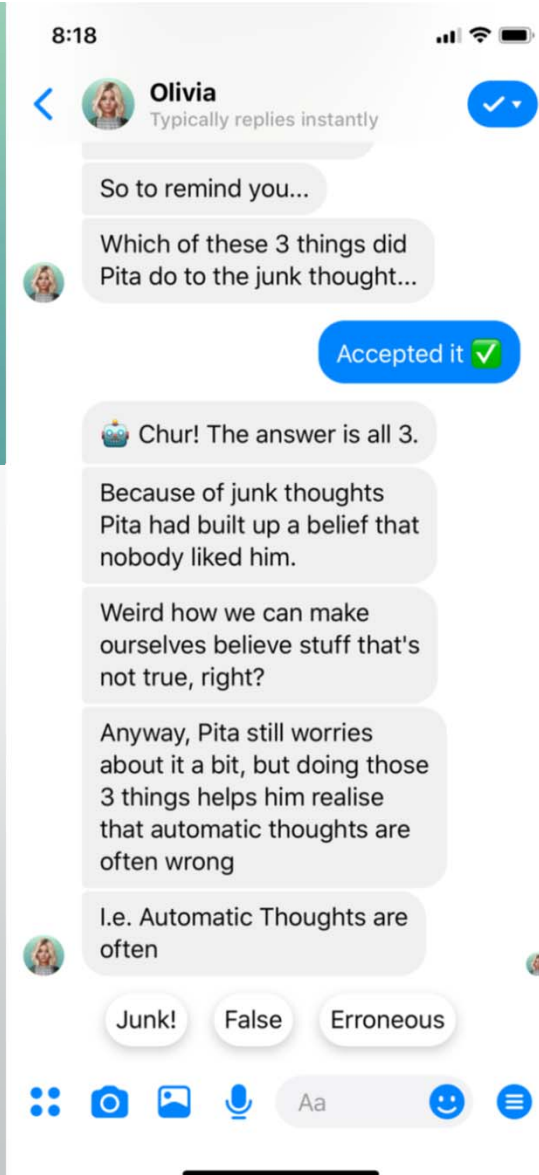
- Extending the HABITs ecosystem with a dialog solution
- CBT and positive psychology
- Relatable avatars / personas
 - But takes care to remind users it's a computer
- 4-week programme
- Delivered through Facebook Messenger
- Back-end content management system supported by Google DialogFlow



What HeadStrong does



- Largely cognitive behavioural therapy (CBT), including
 - Relaxation strategies
 - Problem solving techniques
 - Recognising and tackling negative thoughts
 - Using a gratitude journal
 - Doing positive things (activity scheduling)
- Human (expert counsellor) scripted dialog
 - Easily modified content
 - Ability to amend free-text recognition



Mental health is not an isolated case

- IT in the hands of consumers effective for other health behaviours (things the consumer can change for themselves)
 - Smoking cessation
 - Whittaker R, McRobbie H, Bullen C, Rodgers A, Gu Y. Mobile phone-based interventions for smoking cessation. *Cochrane Database Syst Rev*. 2016 Apr 10;4. (evidence mainly for SMS message interventions)
 - Probably exercise and rehabilitation

www.abc.net.au



Eliza

- Back in 1966, Joseph Weizenbaum made a simple dialogue agent that mimicked a Rogerian psychotherapist

HUMAN: You are like my father in some ways.

ELIZA: What resemblance do you see?

HUMAN: You are not very aggressive but I think you don't want me to notice that.

ELIZA: What makes you think I am not very aggressive?

HUMAN: You don't argue with me.

ELIZA: Why do you think I don't argue with you?

HUMAN: You are afraid of me.

ELIZA: Does it please you to believe I am afraid of you?

HUMAN: My father is afraid of everybody.

ELIZA: What else comes to mind when you think of your father?

HUMAN: Bullies.



Pre-written dialog and affect

- People react emotionally to a voice, or even to text that seems to be coming from an agent: they 'anthropomorphize'
- They don't want to let it down, they look forward to the company*
- Embodiment of the dialog onto a robot presses the point even further!

* B. Kaplan, R. Farzanfar, and R. H. Friedman
Ethnographic interviews to elicit patients' reactions to an intelligent interactive telephone health behavior advisor system. *Proc AMIA Symp.* 1999 : 555–559.



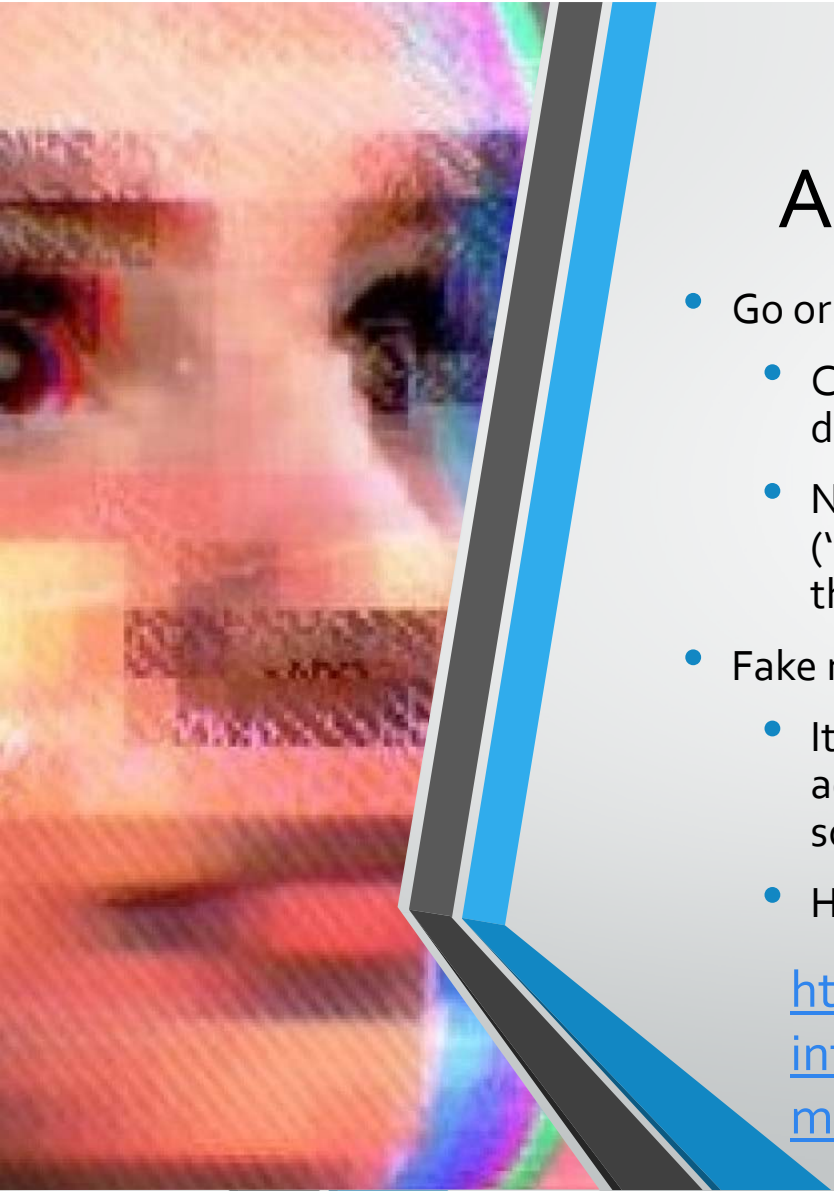
AI dialog challenges and future

- Moving from fully-scripted dialog to personalized and dynamic responses
- Currently investigating transitions into and out of deep-learning enhanced dialog
 - Make system less boring / predictable, build rapport
 - Extend length of engagement, build opportunities to learn things about the user
- Opportunities for deep learning
 - Intent matching (flexibility in how user phrases their input)
 - Dialog policy (how to choose a responding action)
 - End-to-end methods (just machine-learn the most likely / best sequence of words including both patient and counselor)

YOU MUSN'T BE AFRAID
TO DREAM A LITTLE
BIGGER, DARLING.

- Tom Hardy





AI counselling is harder than...

- Go or Chess
 - Counselling is an 'asymmetric game' – patient and counselor are different
 - Need to design two distinct agents to have a reinforcement learning ('self play') loop, and 'winning' is different (and subtle) depending on the patient
- Fake news and Internet trolling
 - It hasn't hurt the (I hope we all agree, **evil**) cause that much if an agent that's dumping abuse on a female Indian politician gets it somewhat wrong
 - High need for coherence and even higher need for getting it right

<https://www.washingtonpost.com/news/the-intersect/wp/2016/03/24/the-internet-turned-tay-microsofts-fun-millennial-ai-bot-into-a-genocidal-maniac/>

My medium-term goal

- Use language encoding and reinforcement learning to...
- Increase 'adherence' to evidence-based e-therapy
 - Make the chatbot more engaging to get significantly more sessions of use so the therapeutic effect (which for now will still be hand crafted) is enhanced
- Give more specific e-therapy
 - Understand what the user wants (or needs, or is indicated for)
 - Deliver the e-therapy with more personalization

Stick

**Stick
with it,
John!**

[Technology](#)[Products](#)[Solutions](#)[News](#)[About](#)[Careers](#)[Contact](#)

Say Hi to Sam and Roman

The future of customer experience is here! See for yourself how Soul Machines™ is putting a human face on AI.

[Talk to one of our Digital People now](#)



<https://www.bbc.com/news/health-50857759>

Will AI take over from humans?

- No. It took humans to design and train the artificial intelligence model.
- This was a research study, and as yet the AI system has not been let loose in the clinic.
- Even when it is, at least one radiologist would remain in charge of diagnosis.
- But AI could largely do away with the need for dual reading of mammograms by two doctors, easing pressure on their workload, say researchers.
- Prof Ara Darzi, report co-author and director of the Cancer Research UK (CRUK) Imperial Centre, told the BBC: "This went far beyond my expectations. It will have a significant impact on improving the quality of reporting, and also free up radiologists to do even more important things."

HABITs Team

prof Sally Merry



Rawiri Wharemate



Dr Grant Christie



Dr Sarah Hopkins



Prof Jim Warren



Dr Terry Fleming



Siu Latu Tony Patolo Stacey Ruru



Tania Cargo



A/Prof Sarah Hetrick



Prof Goodyear-Smith



Dr Hiran Thabrew



Dr Matt Shepherd

