Explainable AI: Beware of Inmates Running the Asylum

Or: How I Learnt to Stop Worrying and Love the Social Sciences
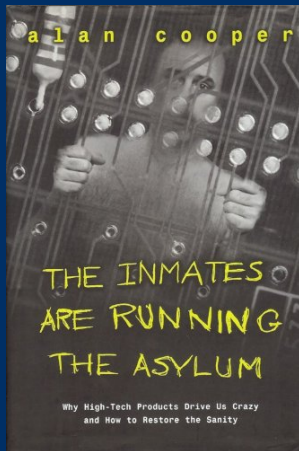
Tim Miller

School of Computing and Information Systems
Co-Director, Centre for AI & Digital Ethics
The University of Melbourne, Australia
tmiller@unimelb.edu.au

THE UNIVERSITY OF
MELBOURNE

29 July, 2021

Alan Cooper (2004): *The Inmates Are Running the Asylum*
Why High-Tech Products Drive Us Crazy and
How We Can Restore the Sanity

Why do we care about explainability?

"...the science team's problems were focused less on exactly what the robot was doing than on why the robot was making particular decisions" — K. Stubbs et al.
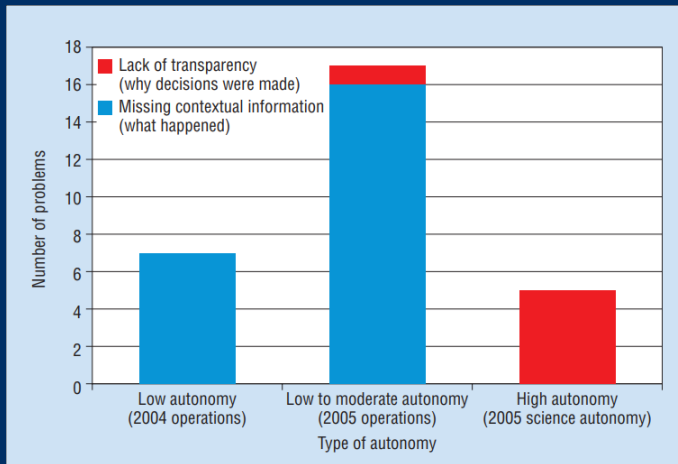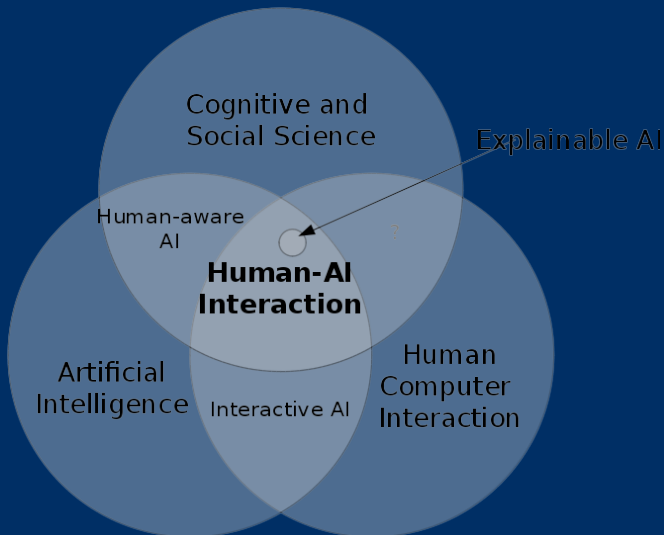
The effects of higher autonomy

Image source: K. Stubbs et al.: Autonomy and Common Ground in Human-Robot Interaction: A Field Study, *IEEE Intelligent Systems*, 22(2):42-50, 2007.

Explainable Artificial Intelligence

Infusing the Social Sciences

A patient has: (1) weight gain; (2) fatigue; and (3) nausea.

GP infers the following most likely causes

Cause	Symptom	Prob.
Stopped Exercising	Weight gain	80%
Mononucleosis	Fatigue	50%
Stomach Virus	Nausea	50%
Pregnancy	Weight gain, fatigue, nausea	15%

Infusing the Social Sciences

A patient has: (1) weight gain; (2) fatigue; and (3) nausea.

GP infers the following most likely causes

Cause	Symptom	Prob.
Stopped Exercising	Weight gain	80%
Mononucleosis	Fatigue	50%
Stomach Virus	Nausea	50%
Pregnancy	Weight gain, fatigue, nausea	15%

The 'Best' Explanation?

A) Stopped exercising and mononucleosis and stomach virus

OR

B) Pregnant

Infusing the Social Sciences

Explanation in artificial intelligence: Insights from the social sciences

Tim Miller

School of Computing and Information Systems, University of Melbourne, Melbourne, Australia

ARTICLE INFO

ABSTRACT

There has been a recent resurgence in the area of explainable artificial intelligence as researchers and practitioners seek to provide more transparency to their algorithms. Much of this research is focused on explicitly explaining decisions or actions to a human observer, and it should not be controversial to say that looking at how humans explain to each other can serve as a useful starting point for explanation in artificial intelligence. However, it is fair to say that most work in explainable artificial intelligence uses only the researchers' intuition of what constitutes a 'good' explanation. There exist vast and valuable bodies of research in philosophy, psychology, and cognitive science of how people define, generate, select, evaluate, and present explanations, which argues that people employ certain cognitive biases and social expectations to the explanation process. This paper argues that the field of explainable artificial intelligence can build on this existing research, and reviews relevant papers from philosophy, cognitive psychology/science, and social psychology, which study these topics. It draws out some important findings, and discusses ways that these can be infused with work on explainable artificial intelligence.

https://arxiv.org/abs/1706.07269

Explanations are *Contrastive*

"The key insight is to recognise that one does not explain events per se, but that one explains why the puzzling event occurred in the target cases but not in some counterfactual contrast case." — D. J. Hilton, *Conversational processes and causal explanation, Psychological Bulletin. 107 (1) (1990) 65–81.*

Why is it a fly?

Type	No. Legs	Stinger	No. Eyes	Compound Eyes	Wings
Spider	8	✘	8	✘	0
Beetle	6	✘	2	✔	2
Bee	6	✔	5	✔	4
Fly	6	✘	5	✔	2

T. Miller. Contrastive Explanation: A Structural-Model Approach, *arXiv preprint arXiv:1811.03163*, 2019. https://arxiv.org/abs/1811.03163

Contrastive Explanation — The Difference Condition

Why is it a fly?

Type	No. Legs	Stinger	No. Eyes	Compound Eyes	Wings
Spider	8	✘	8	✘	0
Beetle	6	✘	2	✔	2
Bee	6	✔	5	✔	4
Fly	6	✘	5	✔	2

T. Miller. Contrastive Explanation: A Structural-Model Approach, *arXiv preprint arXiv:1811.03163*, 2019. https://arxiv.org/abs/1811.03163

Contrastive Explanation — The Difference Condition

Why is it a fly rather than a beetle?

Type	No. Legs	Stinger	No. Eyes	Compound Eyes	Wings
Spider	8	✗	8	✗	0
Beetle	6	✗	2	✔	2
Bee	6	✔	5	✔	4
Fly	6	✗	5	✔	2

T. Miller. Contrastive Explanation: A Structural-Model Approach, *arXiv preprint arXiv:1811.03163*, 2019. https://arxiv.org/abs/1811.03163

Contrastive Explanation — The Difference Condition

Why is it a fly rather than a beetle?

Type	No. Legs	Stinger	No. Eyes	Compound Eyes	Wings
Spider	8	✘	8	✘	0
Beetle	6	✘	**2**	✔	2
Bee	6	✔	5	✔	4
Fly	6	✘	**5**	✔	2

T. Miller. Contrastive Explanation: A Structural-Model Approach, *arXiv preprint arXiv:1811.03163*, 2019. https://arxiv.org/abs/1811.03163

Explanations are *Social*

"Causal explanation is first and foremost a form of social interaction. The verb to explain is a three-place predicate: **Someone** *explains* **something** *to* **someone**. *Causal explanation takes the form of conversation and is thus subject to the rules of conversation."* [Emphasis original]
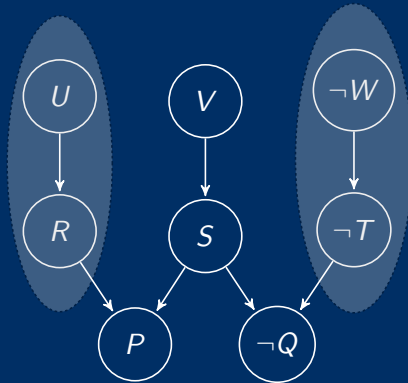
Denis Hilton, Conversational processes and causal explanation, *Psychological Bulletin* 107 (1) (1990) 65–81.

Social Explanation
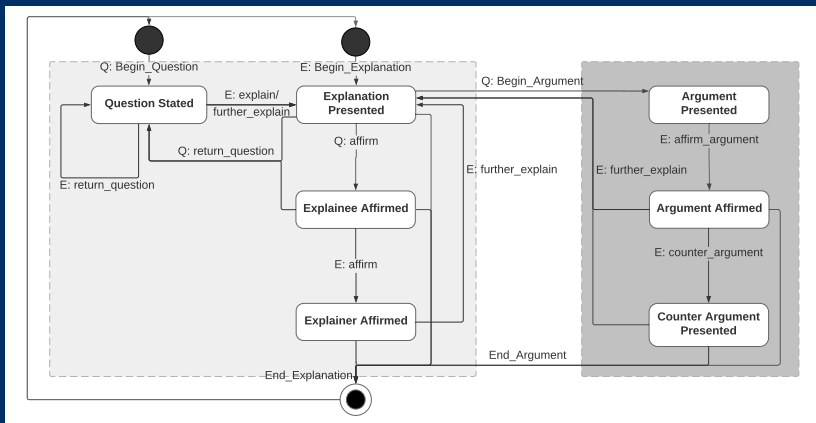
P. Madumal, T. Miller, L. Sonenberg, and F. Vetere. A Grounded Interaction Protocol for Explainable Artificial Intelligence. In *Proceedings of AAMAS 2019*. https://arxiv.org/abs/1903.02409
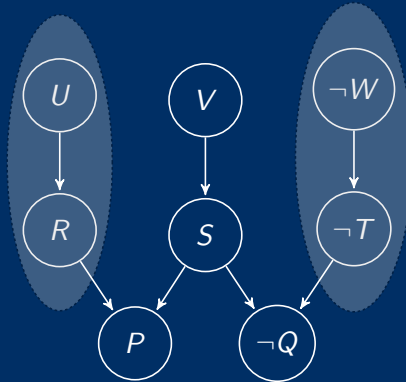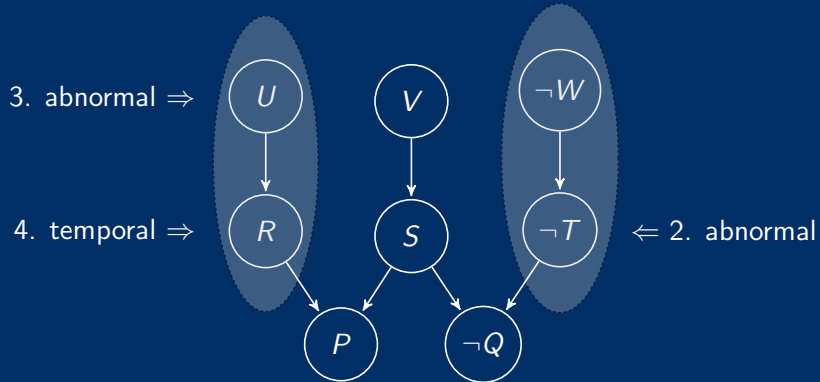
Explanations are *Selected*

4. temporal ⇒

Explanations are *Selected*

Explanations are *Selected*

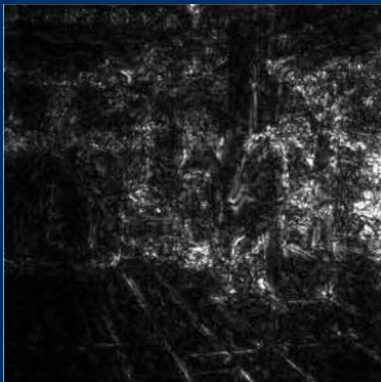3. abnormal $\Rightarrow$ (U) ... (V) ... (¬W) $\Leftarrow$ 1. human intent

4. temporal $\Rightarrow$ (R) ... (S) ... (¬T) $\Leftarrow$ 2. abnormal

(P) ... (¬Q)
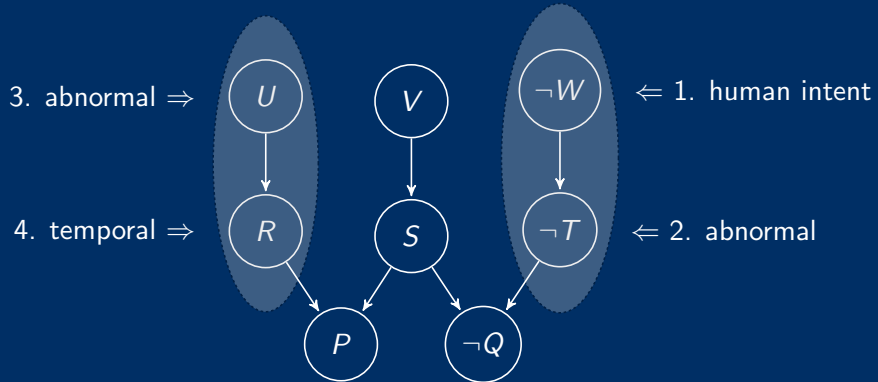
(Not) Infusing Human-Centered Studies

Source: Been Kim: Interpretability – What now? Talk at Google AI.
Saliency map generated using SmoothGrad

Our experience

We have used these insights over a range of techniques:

- Reinforcement learning
- Automated planning
- Computer Vision
- Multi-agent systems

and a range of domains:

- Credit scoring
- Search and rescue
- Illegal fishing
- Game playing

with good results!

Evidence-Driven Models

Generation, selection, and evaluation of explanations is well understood
Social interaction of explanation is reasonably well understood

Evidence-Driven Models

Generation, selection, and evaluation of explanations is well understood
Social interaction of explanation is reasonably well understood

Validation

Validation on human behaviour data is necessary – at some point!

Remember: Hoffman et al., 2018. Metrics for explainable AI: Challenges and prospects. *arXiv preprint arXiv:1812.04608*
https://arxiv.org/abs/1812.04608.

Models

Helping to improve the link between the social sciences and explainable AI.

Wardens, please consider . . .

Models

Helping to improve the link between the social sciences and explainable AI.

Interactions

Helping to study the design of interactions between 'explainable' intelligent agents and people.

Thanks!

Thanks: Prashan Madumal, Piers Howe, Ronal Singh, Liz Sonenberg, Eduardo Velloso, Mor Vered, Frank Vetere, Abeer Alshehri, Ruihan Zhang, Henrietta Lyons, Paul Dourish.

Overview

Explainability is a human-agent interaction problem

The social sciences community perhaps already knows more than the AI community about XAI

Integrating social science research has been useful for my lab:

1. Contrastive explanation
2. Social explanation
3. Selecting explanations

Cross-disciplinary research teams are important!