Digital Ethics – Can We Do Better?

Towards Better Voice Assistants

Saurav Jha, saurav.jha@inria.fr; Inria Nancy, France

Introduction

- Smart (Voice) Assistants: in-home devices connected to the Internet that allow off-site commands customized to a user's need. e.g. Amazon Alexa
- 2019 witnessed 147 million smart assistants sold globally [A]
- Smart assistants scan wake up words
- Everything is being heard
- What does that mean?

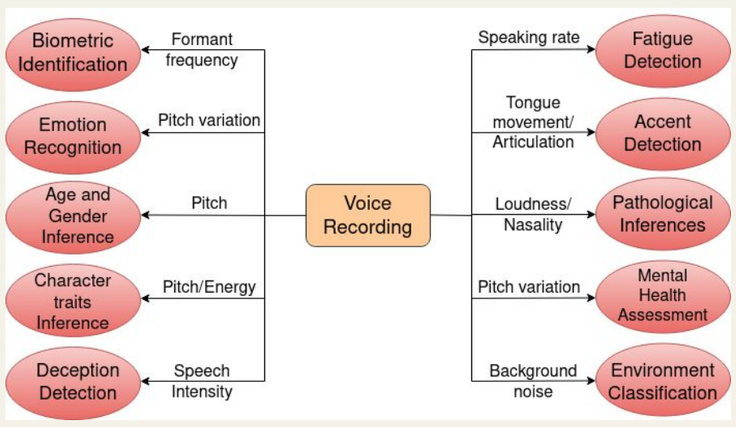

Digital Ethics Issue

➢ What could go wrong?

- Two major issues: Security and Fairness
- Both can be seen to be interlinked in several ways

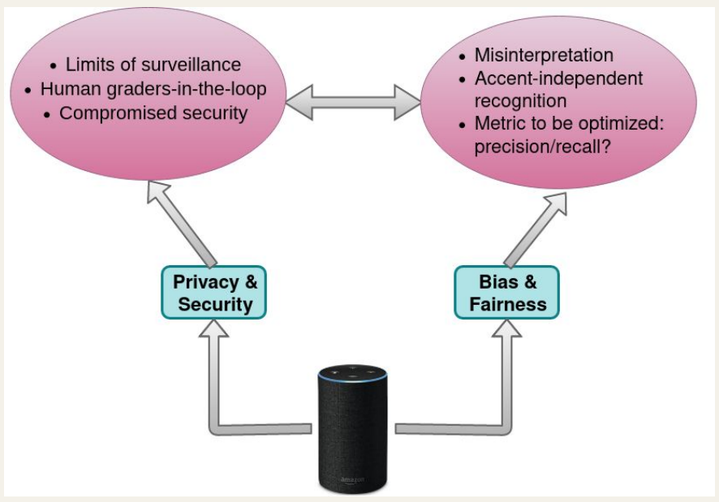

1. Privacy and Security

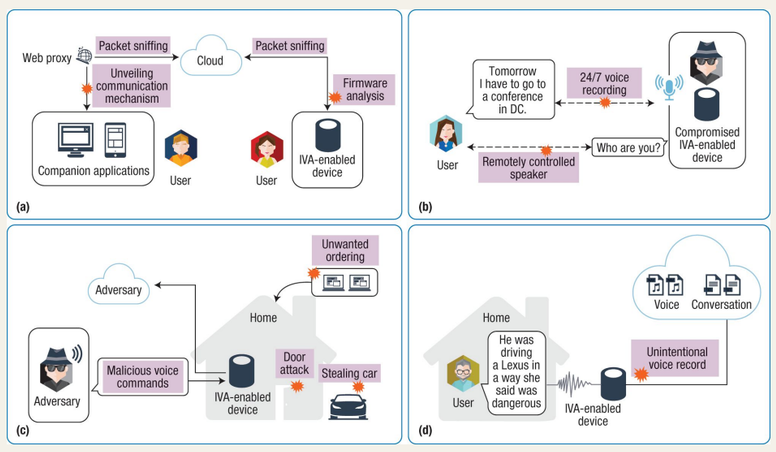

Figure adapted from Chung et al. (2017) showing security vulnerabilities in smart house speakers

2. Bias and Fairness

A more general scenario:

Question: ~~Is a dataset biased?~~
Question: How is a dataset biased?

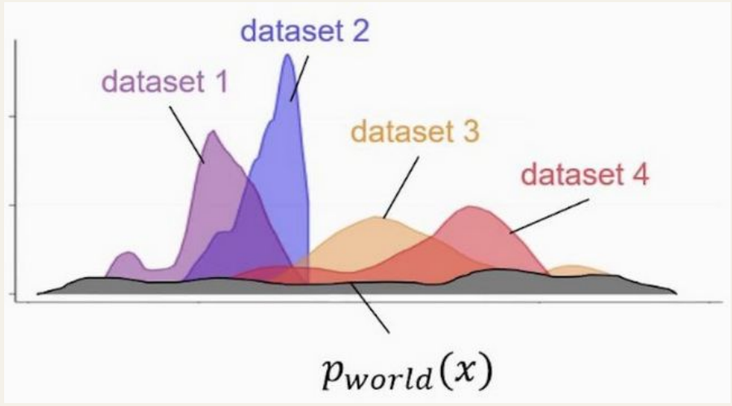

Figure showing the difficulty of sampling i.i.d. from true data distribution $p_{world}(x)$ [2]

- Bias-inducing factors: accent, speed, keywords, colloquial terms, etc.[B]
- Recognition can still be fair to an individual while being unfair to a demographic group and vice-versa.

Discussion & Opinions

- One source of biasness: the data we gather reflects what we choose to look for: out of 177 large US technology companies, % of executives and senior managers being white = 73%, Asian = 21%, Latino = 3%, Black = 1.4%.[3]
- Different value judgements encoded by a speech recognition system may lead to satisfying contradictory fairness properties, i.e., *individual* vs *group fairness* [3]
- Security threats no longer come from hackers and spammers alone, but from powerful nations competing with one another and tech giants with unscrupulous data privacy procedures [4]
- Solving one issue (*security* vs *fairness*) can often help with the other: e.g. Higher precision rates can help tackle vulnerabilities due to unintentional voice records since the device would send data only when it is very confident with speaker recognition
- The pervasive nature of voice assistants can stir an individual moral dilemma. For example, once used to keeping these, can there be a situation where you might be kept from removing these devices due to peer pressure?

Recommendations

- There should be hard policies to guarantee which side to hold responsible for the actions taken by assistants as a result of mistake(s)
- Such policies should be clear at the role of government in seizing data under exceptional circumstances
- As major breakthroughs in AI keep coming from the academia, conferences and journals must emphasize ethical impacts for submissions: e.g. broader impact statements in NeurIPS 2020
- Data-centric AI development should be an equal priority, i.e., paradigm where AI practitioners not just develop code but also data [C]

Bibliography / References

[1] H.Chung, M. Iorga, J. Voas, and S. Lee, "Alexa, can I trust you?" Computer, vol. 50, no. 9, pp. 100-104, 2017.

[2] Kate Saenko, "Is my dataset biased?", ICLR 2021.

[3] Friedler et al., "The (Im)possibility of Fairness", Communications of the ACM, April 2021.

[4] Webroot, "Is Voice Recognition prone to security threats?"

[5] Sinduja Rangarajan, 2018. Silicon valley diversity.

Useful Links

[A] https://tinyurl.com/ps6x5dwn

[B] https://tinyurl.com/2mzkzn2j

[C]https://twitter.com/andrewyng/status/1396922136808202241

Acknowledgments

Special thanks to: Dr. Joseph Timoney, National University of Ireland, Maynooth.