

Q: Let's look at current state of the art for AI.
What types of problems is AI currently good at today,
and why? (with real world examples)


Nisar Ahmed

Smead Aerospace Engineering Sciences

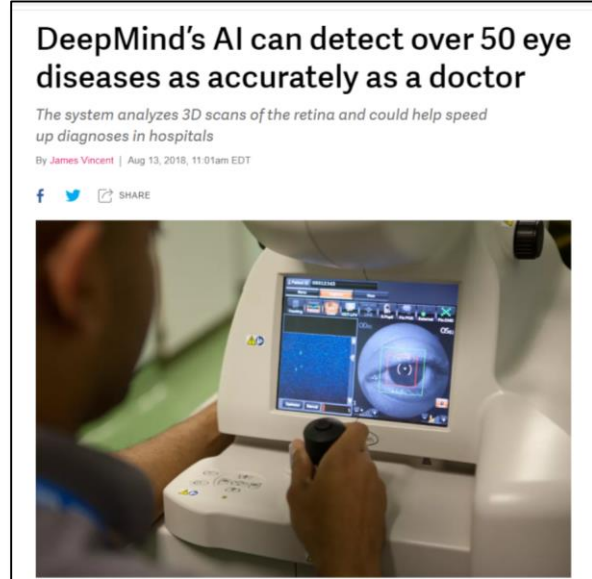
University of Colorado Boulder

“Machine Learning (ML)” : a special domain of AI


- use data to “find algorithm” to turn given list of #'s $x \rightarrow$ desired list of #'s y
- ML methods today: #'s in x and y could represent practically anything...



x: gobs of documents
y: discovery drafts



x: 3D retinal scan data
y: disease diagnosis



x: color images
y: joint movements



x: vehicle sensor data
y: steering + brake commands

How to “find algorithm” from data?

- algorithm = “best fit” often the answer to **some optimization problem**
 - usually: “get predicted y to match desired y (given any x)”
- finding “best fit” for “big problems”: lots of data, computer horsepower
 - Today: storage, CPUs, and GPUs are everywhere & **cheap!**
- closed-worlds, fixed rules, high-fidelity simulations → **ripe for shortcuts!**
 - **Machines “training” against themselves or each other**

