


A major problem with ANNs is that we do not fully understand them.

They obviously work, but:



This stop sign with some carefully added tape patches was interpreted as a speed limit sign of "90" by self driving cars that seems to do everything else right otherwise.

Or



GoogleNet thinks this is a "panda" with less than 58% confidence!

A small perturbation!

GoogleNet thinks this is a "gibbon" with 99.3% confidence!

This means there are risks that we do not understand.

Conclusion: "AI still needs a human "bull shit detector".

Workin in other domains than what we understand as input may make that harder.

Supervised learning requires a supervisor

- An ANN based on supervised learning can not become better than its supervisor.
- If the supervisor is working in the image domain, there is no reason to move the ANN input to a richer domain.