

Apply AI in your field of choice for fun and profit

Lotzi Bölöni

Professor, Computer Science Dept., University of Central Florida

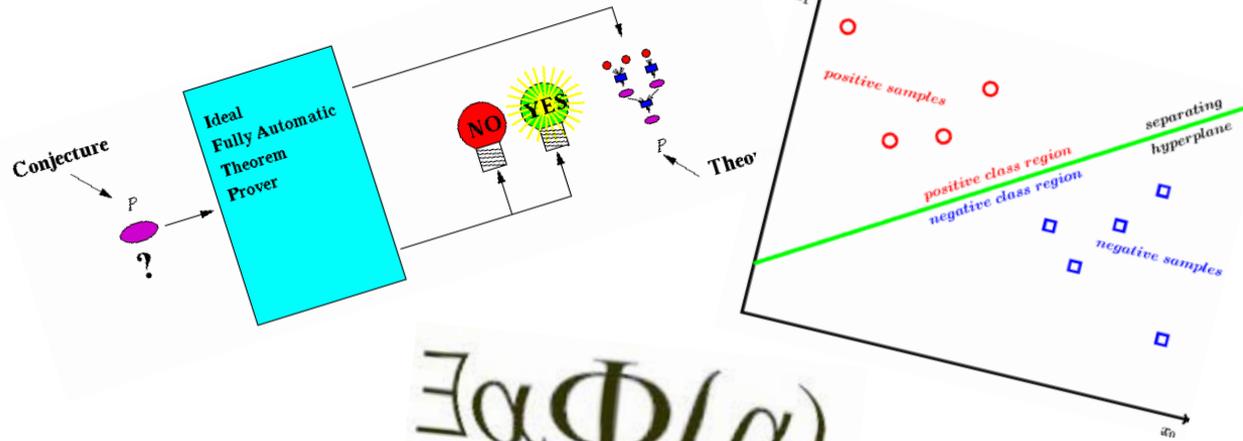
What is this talk about

- This talk captures some of my personal experiences in applying AI in practical fields
 - Working with people who are not in AI
 - Blundering into fields where I am clueless
- It is not supposed to be an exhaustive description of ways to apply AI in other fields

Note: I am going to lump into the AI term all the communities that were historically part of it.

What is AI?

- Theorem proving
- Query answering
- Manipulating symbols
- Estimating probabilities
- Playing games
- Acting in the world
 - Given my knowledge of the world, what action do I need to take in order to further my goals
 - This is the **agent model**



$$\exists \alpha \Phi(\alpha)$$

x $B(x)$ $A(x)$

$\neg K$ $P \rightarrow Q$

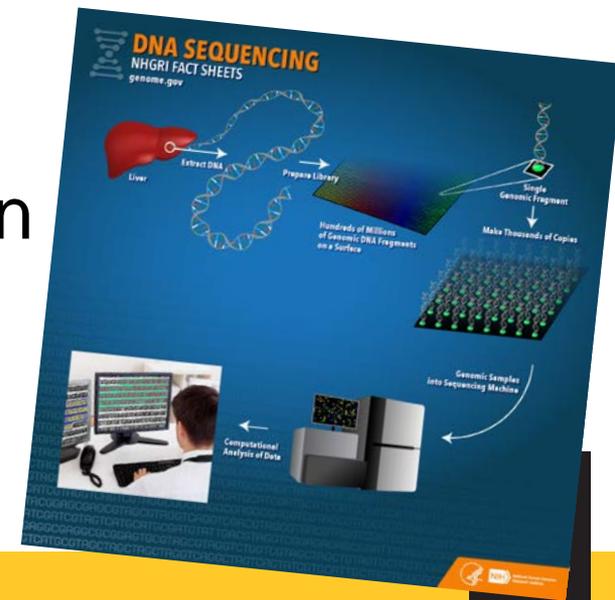
$\forall x$





Why should everyone care about AI

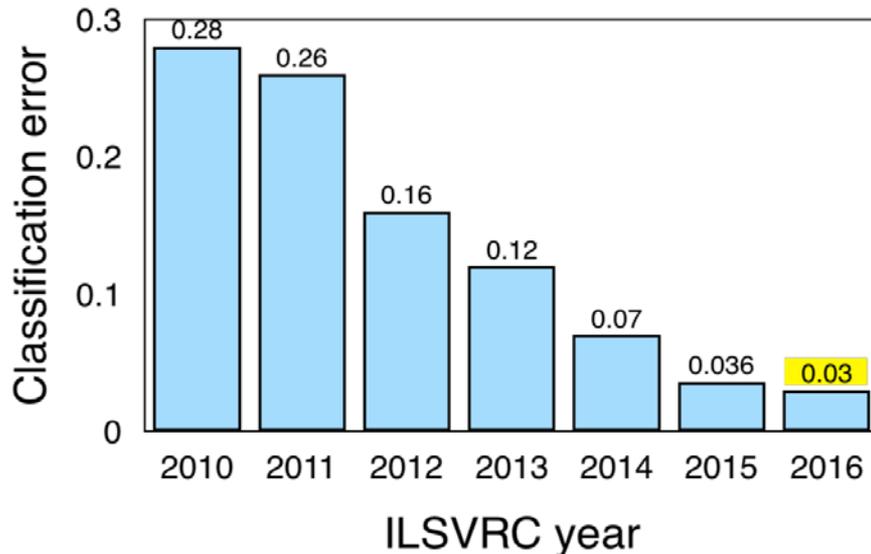
- Let the AI researchers do AI!
 - It is the same as saying that I am in biology, I don't do computers.
- Taking advantage of recent advances in AI allows you to do better whatever you do.
- As AI had made significant leaps in the last five years, you probably want to catch up on the latest results.



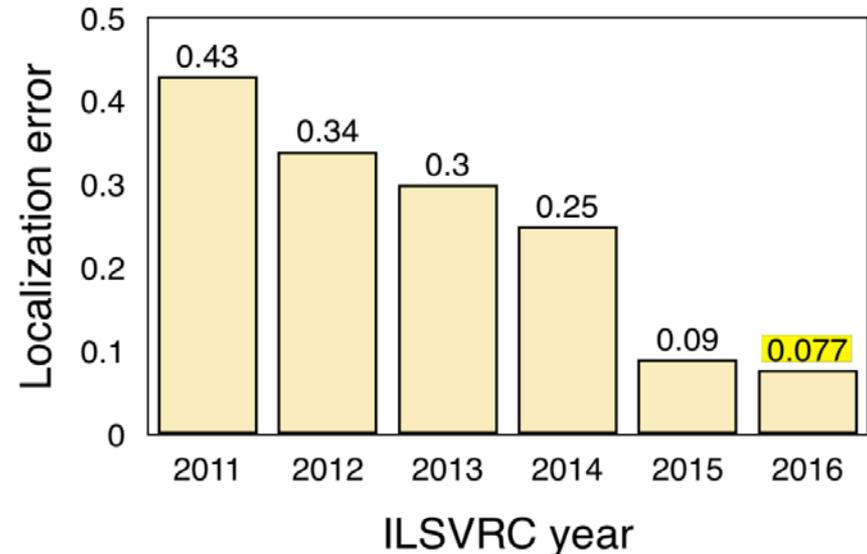
Deep learning changed everything

- So how much performance increase happened?
- ImageNet image processing benchmark
 - Spot the year in which deep learning started working!
 - We are below human error!

Classification



Localization



Laundry list of technologies

- Deep neural networks (even fully connected)
 - For regression or classification
- Deep convolutional neural networks (for vision, video and related tasks) – also in 3D
- Recurrent neural networks (LSTM and GRU) for processing temporal data (or similar)
- Generative adversarial networks (GAN)
- Autoencoders (eg. variational autoencoder VAE)
- Meta-learning, learning-to-learn, one-shot learning

Why should you care about AI, if you are not in AI?

- Amplify the impact of the work you are doing.
 - Better optimization → Better results → Better papers → Fun and profit.
- Focus your human ingenuity on things you really have expertise on
 - ...or care about.
- Let AI take care of the rest

So I am not in AI but I want to take advantage of some of the results

- Some things that worked for me:
 - Take the agent perspective
 - Focus on designing rewards rather than crafting algorithms
 - Consider end-to-end learning (or “learning almost everywhere”)

Take the agent perspective



The agent perspective of AI

- Agent:
 - Sense the world
 - Update your model of the world
 - Decide on the most useful action
 - Act on the world
- Differs from other perspectives on AI because:
 - It is embedded in the world
 - Has a focus on acting (rather than proving, reasoning)

Example from networking

- Consider routing, defined in the traditional way:
 - Incoming packet
 - Look up destination in routing table
 - Send packet on destination port
- Agent perspective
 - Sensing: receive incoming packets (data and control)
 - Update world model based on them
 - Decide whether to transmit, when and what
 - Act (transmit data and control packets)

The decision function can be:

- Hand engineered

- B. Bölöni, D. Turgut, S. Basagni, and C. Petrioli. Scheduling Data Transmissions of Underwater Sensor Nodes for Maximizing Value of Information. In Proc. of IEEE Globecom 2013, pp. 460–465, 2013.

- Use optimization (e.g. mixed integer linear prog.)

- S. Basagni, L. Bölöni, P. Gjanci, C. Petrioli, C.A. Phillips, and D. Turgut. Maximizing the Value of Sensed Information in Underwater Wireless Sensor Networks via an Autonomous Underwater Vehicle. In Proc. IEEE Int'l Conf. on Computer Communications (InfoCom-2014), pp. 988–996, 2014.
- P. Gjanci, C. Petrioli, S. Basagni, C.A. Phillips, L. Bölöni, and D. Turgut. Path Finding for Maximum the Value of Sensed Information in Multi-modal Underwater Wireless Sensor Networks. IEEE Transactions on Mobile Computing, 17:404–418, February 2018.

- Use learning (e.g. RL) etc.

Usefulness of reward function (cont'd)

- It is often more natural to specify a reward than to start engineering behaviors
 - Often, the reward function is given by the application (robot should move to the kitchen...)
 - Often, the application specification omits parts of the reward (...without running over the cat)
- It allows you to specify tradeoffs
 - What happens if not all the goals can be simultaneously satisfied.

Many AI techniques are applied immediately if you have the reward

- Markov Decision Processes, POMDP
 - You need to actually have the reward function
- Reinforcement learning
 - You learn the reward function from interacting with the environment
- Learning from demonstration
 - Reward function is “closeness” to what the human had shown
- Inverse reinforcement learning
 - A variation of learning from demonstration: try to learn the reward function of the human, then use RL.

Sometimes a well-chosen reward can be applied to multiple fields – eg Vol

- Cloud computing

- Bölöni and D. Turgut. Value of information based scheduling of cloud computing resources. *Future Generation Computer Systems*, 71:212–220, June 2017.

- Animal monitoring

- J. Xu, G. Solmaz, R. Rahmatizadeh, L. Bölöni, and D. Turgut. Providing Distribution Estimation for Animal Tracking with Unmanned Aerial Vehicles. To be presented at IEEE GLOBECOM'18, December 2018.

- Internet of Things

- Turgut and L. Bölöni. Value of Information and Cost of Privacy in the Internet of Things. *IEEE Communications Magazine*, 55:62–66, 2017.

- Sensor networks

- F. A. Khan, S. A. Khan, D. Turgut, and L. Bölöni. Optimizing Resurfacing Schedules to Maximize Value of Information in UWSNs. In *Proc of. IEEE Global Communications Conference (GLOBECOM 2016)*, December 2016.

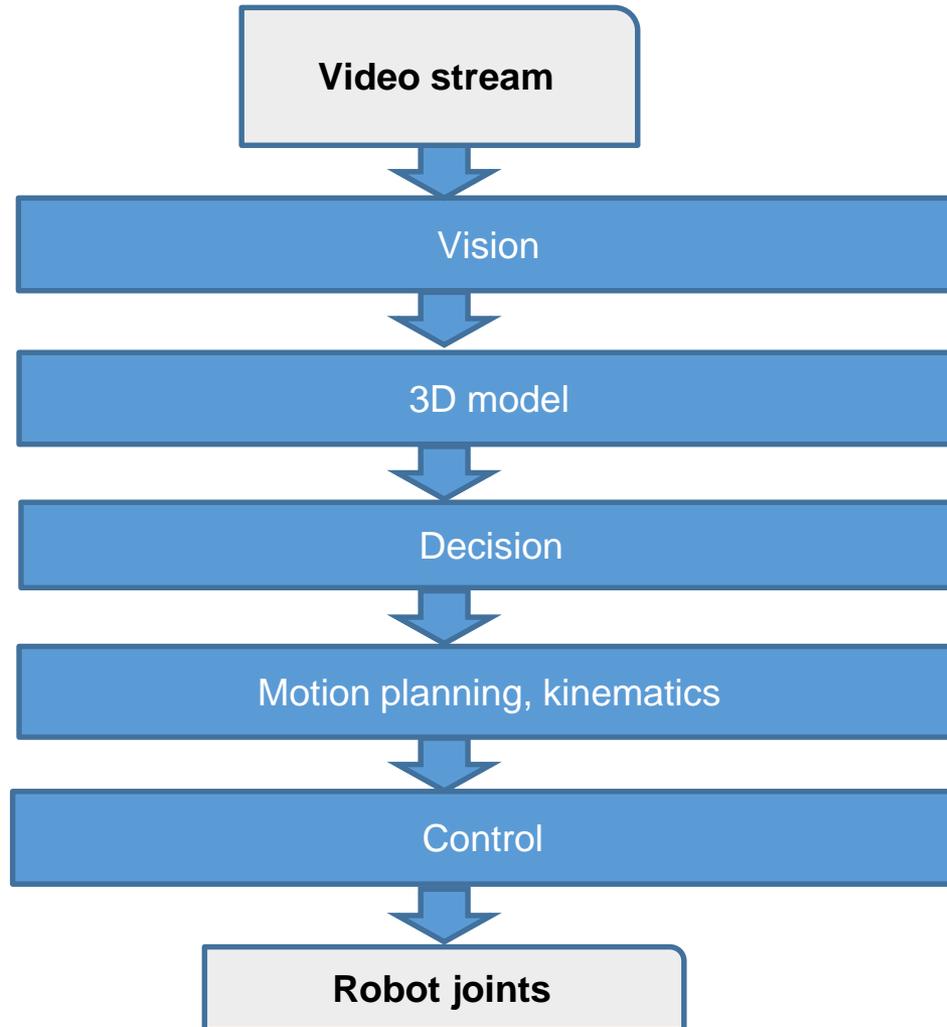
Consider end-to-end learning*

* Or, “learning almost everywhere”

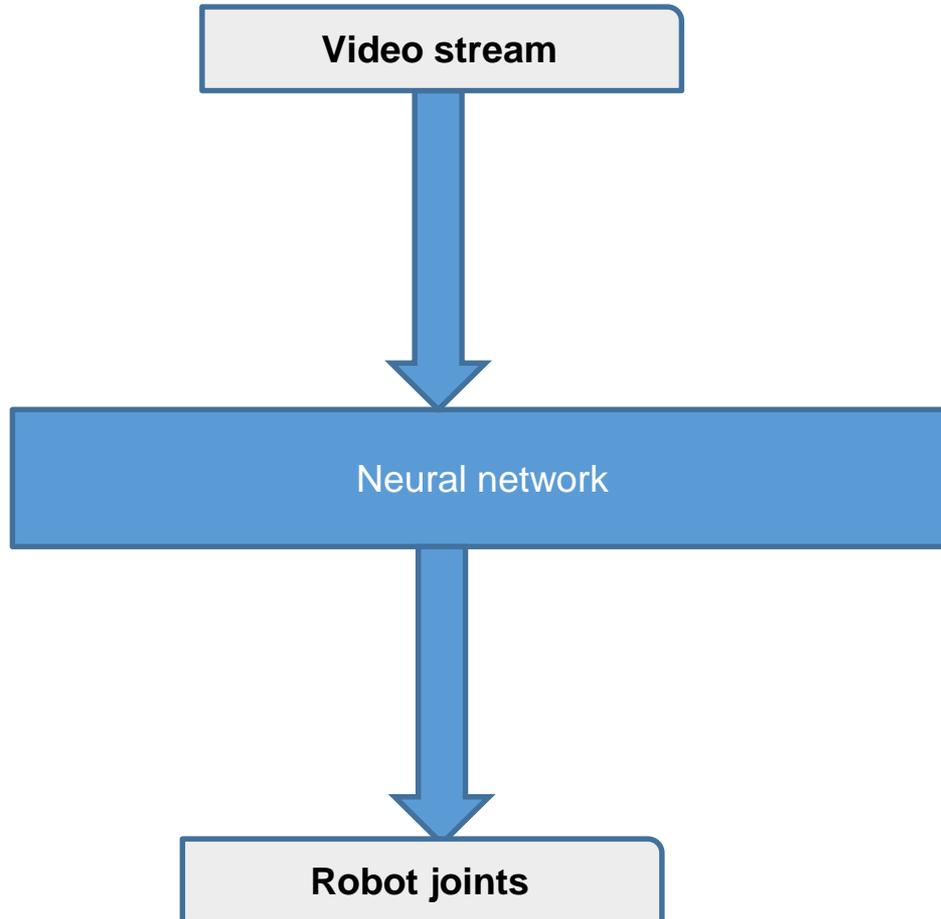
We traditionally create systems by engineering

- Decompose a problem into subproblems
- Then decompose those even further
- Until you reach either a known model, or something that you can solve

Robot control, engineered way



Robot control, end-to-end learning



Use “learning almost everywhere”

- Use engineered solution in places which are
 - Your main contribution
 - In your area of expertise
 - Where you can overperform learning
 - **There are many of these**
- Use learning for things that:
 - You don't know that well
 - Which are likely messy
 - Where learning outperforms handcrafting
 - **... there are also many of these**

Our progression on end-to-end learning for robot manipulator

•Learn manipulation only

- R. Rahmatizadeh, P. Abolghasemi, A. Behal, and L. Bölöni. Learning real manipulation tasks from virtual demonstrations using LSTM and MDN. In Proc. of Thirty-Second AAAI Conf. on Artificial Intelligence (AAAI-2018), February 2018.

- <https://www.youtube.com/watch?v=9vYlIG2ozaM>

•Learn manipulation + vision

- R. Rahmatizadeh, P. Abolghasemi, L. Bölöni, and S. Levine. Vision-Based Multi-Task Manipulation for Inexpensive Robots Using End-To-End Learning from Demonstration. In Proc. of International Conference on Robotics and Automation (ICRA-2018), May 2018.

- <https://www.youtube.com/watch?v=AqQFzoVsJfA>

•Learn manipulation + vision + language + attention

- Abolghasemi, P., Mazaheri, A., Shah, M., & Bölöni, L. (2018). Pay attention!-Robustifying a Deep Visuomotor Policy through Task-Focused Attention. arXiv preprint arXiv:1809.10093.

- https://www.youtube.com/watch?v=xdvNF_R_EkI

Conclusions

- In recent AI had made huge progress in developing new algorithms that can be applied to many problems
- You can amplify the impact of your work (in any field) if you take advantage of these
- I outlined three ways in which this might be done.
- Others exist...

**Thank you
for listening!**

