

Communication-based General Artificial Intelligence

Tomas Mikolov, Facebook

February 2017

The Goal

- AI research is currently very popular...
- But different people see AI as something very different (Image recognition? Machine translation? Data compression?)
- For this talk, we assume “useful AI” is an artificial machine (computer) capable of helping human users in solving wide range of tasks, in a similar way other humans can

Useful AI

- We attempted to identify crucial components of useful AI:
 - Ability to perform tasks for human users
 - Ability to learn and improve
 - Ability to communicate
- We think that removal of any of these requirements would make the machine not useful in a general sense:
 - Remove ability to perform tasks: the machine is not generally useful for us
 - Remove ability to learn: the behavior is fixed and cannot adapt
 - Remove ability to communicate: no way how to specify tasks and obtain results

Additional components

- The previous design can be further extended by adding:
 - Grounding (virtual worlds, 2D / 3D)
 - Multi-modal perception (vision, audio, ...)
 - Communication between AIs
 - ...
- However, we believe that development of even the simplest general AI is very complex and thus start with just the necessary components

The Roadmap to AI

- So we know the goal: we want to develop a machine that can learn to perform novel tasks for us through natural communication - example:

Me: Can you check the weather every day before I go to work to see if it is going to rain, so that I don't forget to bring umbrella?

Machine: But how do I do that?

Me: go to search engine and enter query 'weather new york'
... (some morning week later) ...

Machine: it will rain today!

The Roadmap to AI

- The existing machine learning techniques seem to be insufficient for this goal: deep (recurrent, convolutional) neural networks have excellent performance on supervised tasks, but there is much needed future development of:
 - Unsupervised / reward-based learning
 - Compositional and incremental learning
 - Long term memory
- Currently, there is no existing standard dataset focusing on teaching machines to communicate in natural language while addressing these research problems

CommAI-env

- CommAI-env is an open-source virtual environment focusing on learning to communicate
- Published together with a set of very simple (but still probably unsolvable) communication tasks

CommAI: Evaluating the first steps towards a useful general AI
(Baroni et al, 2017)

<https://github.com/facebookresearch/CommAI-env>

CommAI tasks

- Currently there are several existing datasets, with various degree of complexity

- Example of a basic task:

Teacher: repeat after me: AABBB

Learner: AAB

Teacher: wrong, the correct answer is: AABBB

CommAI tasks

- Currently there are several existing datasets, with various degree of complexity
- Example of a basic task:

Teacher: repeat after me: BBAA

Learner: BBAA

Teacher: good! [reward +1]

Learning to repeat a word (or more generally, a sequence of symbols) is already very challenging for deep neural networks.

CommAI tasks

- The tasks are closely related and share some common structure through natural language

- Example of a task that re-uses previous knowledge:

Teacher: repeat twice after me: ABA

Learner: ABAABA

Teacher: good! [reward +1]

It should be much faster to learn 'repeat twice' task after the Learner already knows how to repeat a sequence once.

CommAI tasks

- Advanced example: learning to count how many red apples and yellow pears I have should be much simpler after the Learner can solve sub-tasks:
 - Understand that objects can have a property like color
 - That red and yellow are names of colors
 - That pears and apples are both countable objects
 - Notion of counting
 - ... and many others
- Even the simplest tasks are actually very challenging: imagine trying to solve such problems in Japanese (or some language you do not know)

Future of general AI research

- We need standard datasets and metrics to compare various attempts to solve communication-based general AI
- CommAI-env is a prototype of communication-based environment, and General AI Challenge from GoodAI is an example of standard dataset defined together with the metrics of success
- The objective function should reflect the learning speed: we aim to build Learners that can learn from as few examples as possible to perform novel tasks

Open research problems

Unsupervised / reward-based learning:

- How can the machine “learn to learn”: knowledgeable Learner should modify its own behavior even when no explicit reward signal is present
- Learner should be able to memorize new facts and abilities without explicit instructions

Open research problems

Compositional and incremental learning:

- How can the Learner build new skills by re-using the existing skills, ie. without learning from scratch solution to every new problem?

Open research problems

Long term memory:

- What should be the structure of the long term memory? How should it be updated?
- All these research problems are probably very related. One can probably not solve compositional and incremental learning without having any way to form persistent long term memory.

Conclusion

- We may want to be more goal-oriented when we talk about strong or general AI
- Communication seems to be necessary for useful strong AI
- To achieve progress, researchers need to have standard datasets to compare various approaches, and incentives to work on very difficult unsolved problems