A Roadmap Towards Machine Intelligence

Some (Not So) Recent Papers by Tomáš Mikolov

Vítek Novotný, 2021-04-22

Tomáš Mikolov (* 1982 Šumperk)

- 2007: [Language Modeling for Spech Recognition in Czech]() (Mgr. thesis, BUT)
- 2010: Stays at Johns Hopkins University (F. Jelínek) and Montreal (Bengio)
- 2010–2014: Google Brain, 2014–2020: Facebook AI Research
- 2020–: Czech Institute of Informatics, Robotics and Cybernetics (CIIRC CTU)

A Roadmap Towards Machine Intelligence (2018) I

1. **Introduction**

   “Given the current availability of powerful hardware and large amounts of machine-readable data, as well as the widespread interest in sophisticated machine learning methods, the times should be ripe for the development of intelligent machines.”

2. **Desiderata for an intelligent machine**
   1. **Ability to communicate**
   2. **Ability to learn**

A Roadmap Towards Machine Intelligence (2018) II

3. **A simulated ecosystem to educate communication-based intelligent machines**
    1. **High-level description of the ecosystem**
        - **Agents:** Teacher (and Reward) ↔ Learner ↔ Environment
        - **Interface channels:** Input and output bitsteams
        - **Reward, Incremental structure, Time off, Evaluation**
    2. **Early stages of the simulation**
        - **Preliminaries, Notation**

(continued on the next slide)

The Learner learns to issue Environment commands

Input:	*Output:*
T: give order **@E:** I move.	
	fdakjf13ktj53u9gs5 5 3#.
character soup without addressee, no reward	
	...
	@T: dsfjkasjfk;aljsf.
character soup with wrong addressee, no reward	
	...
	@E: fglk4$3wfgg.
right addressee, character soup, no reward	
	...
	give order **@E:** I move.
E: you moved.	
R: 1.	
Learner moves and is rewarded, although the solution is inefficient	
T: give order **@E:** I move.	
	...
	@E: I move.
E: you moved.	
R: 1.	
Learner figured out minimal string leading to reward	

Associating language to actions

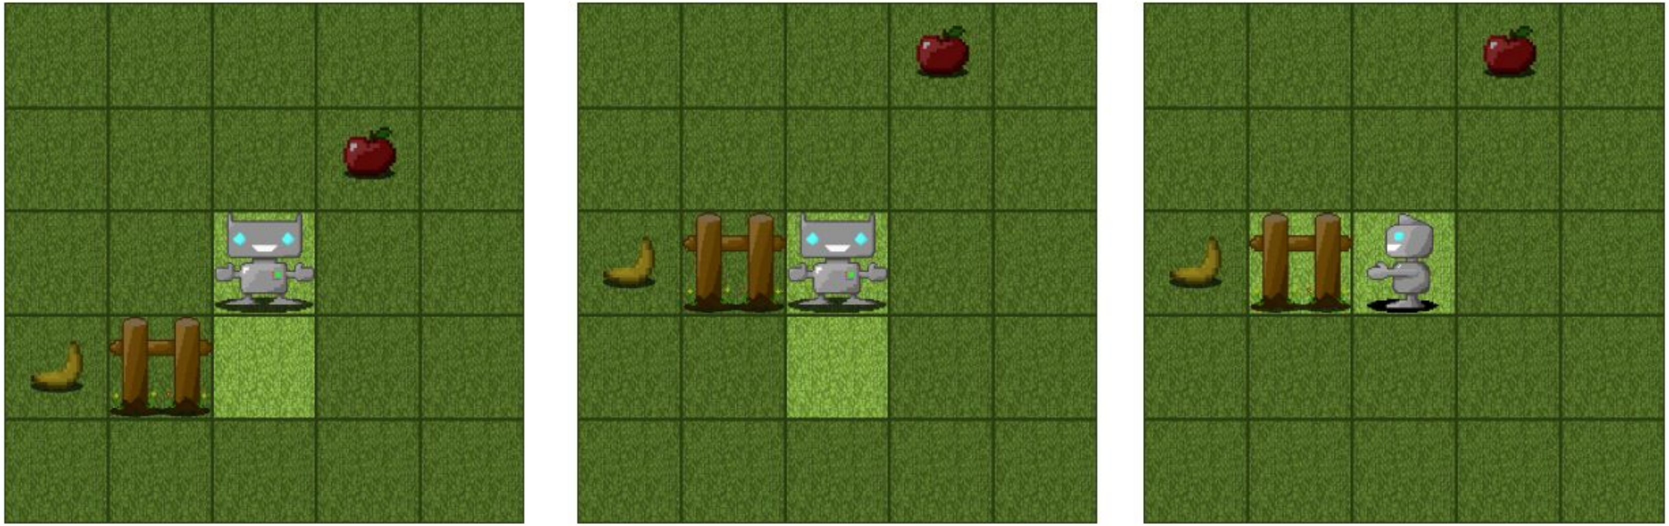

Input: *Output:*
T: move, turn right and move.
@E: I move.

Input: *Output:*
E: you moved.
@E: I turn right.

Input: *Output:*
E: you turned right.
@E: I move.
E: you can't move.

Figure 2: The Teacher gives the order *move, turn right and move*. The Learner is however blocked by a barrier after it turns right, and reward is delayed, because the order is not executed. (Best viewed in color.)

Learning to generalize I

Input:	*Output:*
T: turn and move.	
	@E: I turn.
E: *(no action)*	
	...
	@E: I turn left.
E: you turned left.	
	@E: I move.
E: you can't move. there is a barrier in front of Learner	
	...
	@E: I turn right.
E: you turned right.	
	@E: I move.
E: you moved. **R:** 1.	

Learning to generalize II

Input:	*Output:*
T: pick an object.	
	@E: I pick an object.
E: *(no response)*	
	@E: I pick the apple.
E: *(no response)*	
	@E: I pick the pear.
E: you picked the pear. **R:** 1.	

As we just illustrated, initially the Learner will apply an exhaustive search strategy, listing all objects it knows of to find one it can pick. Later, the Teacher should teach the *look* command, and the Learner should be able to discover a faster strategy than exhaustive search:

Input:	*Output:*
T: pick an object.	
	@E: I look.
E: you see a pear.	
	@E: I pick the pear.
E: you picked the pear. **R:** 1.	

Learning to generalize III

Input:

T: move and look.

Output:

@E: I move.

E: you moved.

@E: I look.

E: you see grass.
T: move and look.

@E: I move.

E: you moved.

@E: I look.

E: you see an apple.
R: 1.
T: this is called find an apple.

Interactive communication

Input:	*Output:*
	@T: how to find an apple.
T: turn right and move and move and pick the apple.	
	@E: I turn right.
E: you turned right.	
	@E: I move.
E: you moved.	
	@E: I move.
E: you moved.	
	@E: I pick the apple.
E: you picked the apple. **R:** 1.	

- **Algorithmic knowledge**

  “Thus, we believe that successful construction of intelligent machinescould automate computer programming, which will likely be done in the future simply through communication in natural language.”

Interacting with the trained intelligent machine

Input:	*Output:*
Bob: I just spoke to the doctor, who said my mother needs to move for at least one hour per day, please make sure she does get enough exercise.	
. . .	
following conversation takes place the next day:	
	@Alice: Are you in the mood for some light physical exercise today?
Alice: Yes, but no more than 15 minutes, please.	
	@INTERNET: [query search engine for keywords *elderly, light activity, 15 minutes*]
shortly afterwards. . .	
	@Alice: I have downloaded a YouTube video with a 15-minute yoga routine, please watch it whenever you can.

A Roadmap Towards Machine Intelligence (2018) X

4. **Towards the development of intelligent machines**
    1. **Types of learning**
    2. **Long-term memory and compositional learning skills**
    3. **Computational properties of intelligent machines**

    “Since there are many Turing-complete computational systems, one may wonder which one should be preferred as the basis for machine intelligence. We cannot answer this question yet, however we hypothesize that the most natural choice would be a system that performs computation in a parallel way, using elementary units that can grow in number based on the task at hand. The growing property is necessary to support the long-term memory, if we assume that the basic units themselves are finite. An example of an existing computational system with many of the desired properties is the cellular automaton of Von Neumann et al. (1966). [...]”

Possible Computational Systems for the Intelligent Machine

- KRUSZEWSKI, Germán; MIKOLOV, Tomas. Combinatory Chemistry: Towards a Simple Model of Emergent Evolution. In: Artificial Life Conference Proceedings. One Rogers Street, Cambridge, MA 02142-1209 USA journals-info@ mit. edu: MIT Press, 2020. p. 411-419. Available at DOI: 10.1162/isal_a_00258
- CISNEROS, Hugo; SIVIC, Josef; MIKOLOV, Tomas. Visualizing computation in large-scale cellular automata. In: Artificial Life Conference Proceedings. One Rogers Street, Cambridge, MA 02142-1209 USA journals-info@ mit. edu: MIT Press, 2020. p. 239-247. Available at DOI: 10.1162/isal_a_00277

Combinatory Chemistry: Towards a Simple Model of Emergent Evolution (2020)

- Introduces **combinatory chemistry**: an artificial algorithmic chemistry based on combinatory logic (CL), a minimalistic computational system that was independently invented by Schonfinkel, Von Neumann, and Haskell Curry.
- A system consists of CL expressions that react following reduction rules, plus random condensation and cleavages.
- They apply a heuristic approach, which emulate the effects of having a much larger system, and study the emergence of complex structures.
- They seek to use combinatory chemistry for explaining the emergence of evolvability, one of the central questions in Artificial Life.

Visualizing Computation in Large-Scale Cellular Automata

- Large-scale cellular automata (CA) have high time and memory complexity.
- They propose several coarse-graining approaches for cellular automata, which approximate the original cellular automata by merging cells with losing as few information about the states as possible.
- This allows them to visualize large-scale CAs and see emergent structures.

“Viewing space-time diagrams of cellular automata is akin to visualizing a foreign computer design. Cellular automata are manipulating information, registers and instructions in parallel in the form of cell states. [...] Future work could focus on identifying some known simple computational primitives within cellular automata and understanding how our visualization can help to find them.”