Certificate I: Understanding AI and Machine Learning in Africa

Course AIMLO1: Artificial Intelligence – Past, Present, and Future

Module 4: Future Challenges

Lecture 2: Self-programming and Self-learning Systems

Learning Objectives

- 1. Explain the difference between a self-programming system and a self-learning system.
- 2. Provide examples of self-programming systems
- 3. Provide examples of self-learning systems

Lecture Contents

- 1. The quest for self-programming and self-learning computer systems
- 2. Self-programming based on deep learning and natural language processing (NLP)
- 3. Self-learning in developmental robotics
- 4. Reinforcement learning and self-supervised learning
- 5. Lecture summary
- 6. Recommended reading & references

Carnegie Mellon University Africa

 The automatic generation of computer programs has been one of the main challenges of AI since the beginning

Also known as program synthesis and self-programming

- The first symbolic GOFAI approaches to AI aimed to use general-purpose knowledge to
 - Generate new text
 - Solve mathematical problems
 - Create new computer programs

It is considered the first useful AI program. It was intended to solve many different problems, rather than one specific problem, using the same reasoning mechanism for each problem.

It eventually evolved into the Soar cognitive architecture.

• For example, the General Problem Solver (Newell et al. 1959)

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future Carnegie Mellon University Africa

More recently, AI has also focused on self-learning system

That can learn with no or minimal supervision from humans

We say more about this approach later in this lecture

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future Carnegie Mellon University Africa

Self-Programming

- The self-programming system challenge has recently received a significant boost through the combination of deep learning methods for NLP
- DeepCode (Balog et al. 2016) is a code generator that uses a neural network
 - To predict the properties of the program
 - That can produce the required outputs
 - Given specific the inputs
 - This neural network is used to augment search-based techniques to solve the inductive program synthesis problem for simple types of programs typical of programming competitions

- SketchAdapt (Nye et al., 2019) is a system that learns, without direct supervision
 - When to rely dynamically on pattern recognition
 - When to perform symbolic search for explicit reasoning
- This mimics the human ability to dynamically incorporate pattern recognition from examples and reasoning to solve programming problems from examples or a specification expressed in natural language

Self-Programming

• GPT-3 General Pre-trained Transformer (Brown et al. 2020) from OpenAl has potential for automatic program synthesis

We met GPT-3 briefly in Lecture AIML01-02-02.

- Advanced large-scale Transformer neural network
- 175 billion parameters
- Generates natural language text
- It can generate new text
 - Without the need of further training
 - Without task-specific fine tuning of its parameters

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future Carnegie Mellon University Africa

- GPT-3 has been evaluated with few-shot demonstrations
 - Given the task description
 - Given one or few examples
- It can also perform several tasks that require on-the-fly reasoning
 - Unscrambling words
 - Using a novel word in a sentence
 - Performing arithmetic

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future Carnegie Mellon University Africa

- GPT-3 can produce samples of news articles which human evaluators have difficulty distinguishing from articles written by humans
- GPT-3 has been used for generating programs, such as for the code to create the Google home page (Bussler 2020, Heaven 2000)

Attempts to design AI systems and robots that autonomously learn without supervision from humans have recently been realized in developmental robotics

- Taking inspiration from child development —
- To design robots that go through stages of development
- For the incremental acquisition of sensorimotor and cognitive skills (Cangelosi & Schlesinger 2015)



https://mitpress.mit.edu/books/developmental-robotics

Module 4: Future Challenges Lecture 2: Self-programming Systems; Slide 11

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future

Developmental robots use intrinsic motivation mechanisms

- Implemented with reinforcement learning, for example
- To allow them to initiate and manage self-learning
- via curiosity-driven mechanisms
- for open-ended, cumulative acquisition of skills



https://mitpress.mit.edu/books/developmental-robotics

Module 4: Future Challenges Lecture 2: Self-programming Systems; Slide 12

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future

Another example of self-learning AI is the AlphaGo Zero system

we met AlphaGo Zero in Lecture AIML01-03-04

Here artificial agents play the game Go against each other, bootstrapping their final learning capabilities

This led to the acquisition of skills that far outperformed the skills of the best human players and previous versions of AlphaGo (Silver et al., 2017)



https://www.bbc.com/news/technology-35785875

Carnegie Mellon University Africa

AlphaGo Zero achieved this based purely on reinforcement learning without any prior supervised training

AlphaGo uncovered several innovative strategies that greatly surprised expert players,

This demonstrates the potential for AI to augment human abilities and exceed human performance



AlphaGo - The Movie

https://www.youtube.com/watch?v=WXuK6gekU1Y

Carnegie Mellon University Africa





A variant of unsupervised learning where the data provides the supervision

Also referred to as predictive learning

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future

Carnegie Mellon University Africa



Self-supervised – predictive – learning leverages the power of unsupervised methods, such as autoencoders and word embeddings, for automatically extracting partial information from noisy or incomplete input data to predict the rest of the data.

Certificate I: Understanding AI and Machine Learning in Africa Course AIMLO1: Artificial Intelligence – Past, Present, and Future

Carnegie Mellon University Africa

Lecture Summary

- 1. Self-programming systems focus on automatic generation of computer programs
- 2. Self-programming systems also focus on using general purpose knowledge and reasoning to solve many different problems
- 3. Today, self-programming systems are levering recent advances in natural language processing (NLP)
- 4. Self-learning systems can learn with little or no supervision from humans
- 5. Artificial curiosity drives self-learning in developmental robotics
- 6. Self-learning systems can also use reinforcement and self-supervised learning

Recommended Reading

Peng T. (2019) LeCun Cake Analogy 2.0. https://medium.com/syncedreview/yann-lecun-cake-analogy-2-0-a361da560dae

Carnegie Mellon University Africa

References

Balog, M., Gaunt, A. L., Brockschmidt, M., Nowozin, S., & Tarlow, D. (2016). DeepCoder: Learning to write programs. arXiv preprint arXiv:1611.01989.

https://www.researchgate.net/publication/309738607_DeepCoder_Learning_to_Write_Programs

Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler DM, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D (2020) Language models are few-shot learners. arXiv:2005.14165. https://arxiv.org/abs/2005.14165

Bussler, F. (2020) Will the latest AI kill coding? https://towardsdatascience.com/will-gpt-3-kill-coding-630e4518c04d

- Cangelosi, A., & Schlesinger, M. (2015). Developmental Robotics: From Babies to Robots. MIT press. https://mitpress.mit.edu/books/developmental-robotics
- Heaven, W. D. (2020). OpenAl's new language generator GPT-3 is shockingly good—and completely mindless. MIT Technology Review.

https://www.technologyreview.com/2020/07/20/1005454/openai-machine-learning-language-generator-gpt-3-nlp/

References

Nye, M., Hewitt, L., Tenenbaum, J., Solar-Lezama, A. (2019). Learning to Infer Program Sketches, Proceedings of the 36th International Conference on Machine Learning, PMLR 97:4861-4870, 2019. http://proceedings.mlr.press/v97/nye19a.html

Newell, A., Shaw, J.C., Simon, H.A. (1959). Report on a general problem-solving program. Proceedings of the International Conference on Information Processing. pp. 256–264. http://bitsavers.informatik.uni-stuttgart.de/pdf/rand/ipl/P-1584_Report_On_A_General_Problem-Solving_Program_Feb59.pdf

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of Go without human knowledge. Nature, 550(7676):354. https://www.researchgate.net/publication/320473480_Mastering_the_game_of_Go_without_human_knowledge

Carnegie Mellon University Africa