DEEP LEARNING ALGORITHM USED IN ROBOTICS

Ms.M.Prithi
Assistant Professor,Department of Computer Applications
Marudhar Kesari Jain College for women, Vaniyambadi, Vellore.

Ms,R.Sandhiya
Assistant Professor,Department of Computer Applications
Marudhar Kesari Jain College for women, Vaniyambadi, Vellore.

Abstract

Advances in deep learning over the last decade have led to a flurry of research in the application of deep artificial neural networks to robotic systems, with at least thirty papers published on the subject between 2014 and the present. This review discusses the applications, benefits, and limitations of deep learning vis-à-vis physical robotic systems, using contemporary research as exemplars. It is intended to communicate recent advances to the wider robotics community and inspire additional interest in and application of deep learning in robotics

Deep learning allows robots to do better analysis and detection of in vulgar things through more detailed and complex data.

Here are some examples of applications of deep learning:

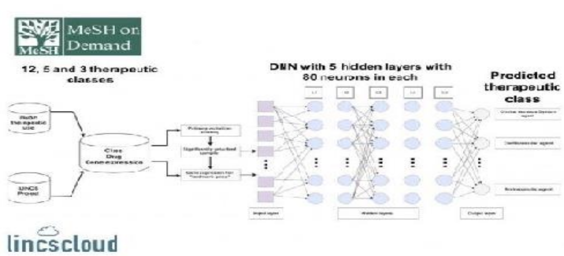

Robots use deep networks that are used to categorize drugs and process large amounts of cells and even to identify biomaterial from the blood.

- **Physics of particle**

Robots with deep learning systems are being used used to detect the particles like the Higs Boson by detecting the reactions and defining which particle is now there.

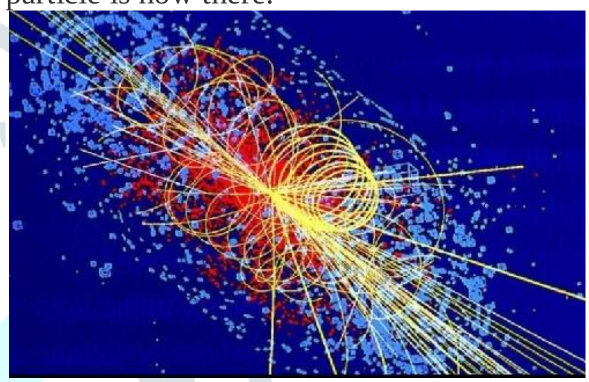

How is deep learning used in robotics?

There are lots of possible applications of deep learning concepts in Robotics.One of the most popular applications is Sentiment Analysis of Images being processed in real time.However, much of such endeavors are limited to researches and hobby stuff as performing deep learning algorithms on a scaled robotic platform (that is recording lots of abstract images and video content) is in itself a challenge.

A big chunk of it can be credited to

- Heavy processing power
- Scaled Cost
- Higher Processing time
- Inefficient Image Processing Algorithms

A big leap in the field of Artificial Intelligence is marked by efficient sentiment analysis. It generally involves semantic context mining and related semantic data mining from objects (movable or immovable), scenes or images and generating tags and metadata using supervised or unsupervised machine learning models.

Deep learning high level trajectory of deep learning with robotics

Ultimately, the underlying philosophy that prevails in the deep learning community is that every part of a complex system can be made to "learn." Thus, the real power of deep learning does not come from using just one of the structures described in the previous section as a component in a robotics system, but in connecting parts of all of these structures together to form a full system that learns throughout. This is where the "deep" in deep learning begins to make its impact – when each part of a system is capable of learning, the system as a whole can adapt in sophisticated ways. Neuroscientists are even starting to recognize that many of the patterns evolving within the deep learning community and throughout artificial intelligence are starting to mirror some of those that have previously evolved in the brain . Doya identified that supervised learning methods (Structures A and C) mirror the function of the cerebellum, unsupervised methods (Structure B) learn in a manner comparable to that of the cerebral cortex, and reinforcement learning is analogous with the basal ganglia . Thus, the current trajectory of advancement strongly suggests that control of robots is leading toward full cognitive architectures that divide coordination tasks in a manner increasingly analogous with the brain .

Deep learning in robotics

The robotics community has identified numerous goal for robotics in the next 5 to 20 years. These include, but certainly are not limited to, human-like walking and running, teaching by demonstration, mobile navigation in pedestrian environments, collaborative automation, automated bin/shelf picking, automated combat recovery, and automated aircraft inspection and maintenance, and robotic disaster mitigation and recovery.

This paper identifies seven general challenges for robotics that are critical for few reaching these goals and for which DNN technology has high potential for impact:

Challenge 1: ***Learning complex, high-dimensional, and novel dynamics***.

Analytic derivation of complex dynamics requires human experts, is time consuming, and poses a trade-off between state dimensionality and tractability. Making such models robust to uncertainty is difficult, and full state information is often unknown. Systems that can quickly and autonomously adapt to novel dynamics are needed to solve problems such as grasping new objects, traveling over surfaces with unknown or uncertain properties, managing interactions between a new tool and/or environment, or adapting to degradation and/or failure of robot subsystems. Also needed are methods to accomplish this for systems that possess hundreds (or even thousands) of degrees of freedom, exhibit high levels of uncertainty, and for which only partial state information is available.

Challenge 2: ***Learning control policies in dynamic environments***.

As with dynamics, control systems that accommodate high degrees of freedom for applications such as multi-arm mobile manipulators, anthropomorphic hands, and swarm robotics are needed. Such systems will be called upon to function reliably and safely in environments with high uncertainty and limited state information.

Challenge 3: ***Advanced manipulation***.

Despite advances achieved over 3 decades of active research, robust and general solutions for tasks such as grasping deformable and/or complex geometries, using tools, and actuating systems in the environment (turn a valve handle, open a door, and so forth) remain elusive – especially in novel situations. This challenge includes kinematic, kinetic, and grasp planning inherent in tasks such as these.

Challenge 4: ***Advanced object recognition***.

DNNs have already proven to be highly adept at recognizing and classifying objects . Advanced application examples include recognizing deformable objects and estimating their state and pose for grasping, semantic task and path specification (e.g., go around the table, to the car, and open the trunk), and recognizing the properties of objects and surfaces such as sharp objects that could pose a danger to human collaborators or wet/slippery floors.

Challenge 5: *Interpreting and anticipating human actions.*

This challenge is critical if robots are to work with or amongst people in applications such as collaborative robotics for manufacturing, eldercare, autonomous vehicles operating on public thoroughfares, or navigating pedestrian environments. It will enable teaching by demonstration, which will in turn facilitate task specification by individuals without expertise in robotics or programming. This challenge may also be extended to perceiving human needs and anticipating when robotic intervention is appropriate.

Challenge 6: *Sensor fusion & dimensionality reduction.*

The proliferation of low-cost sensing technologies has been a boon for robotics, providing a plethora of potentially rich, high-dimensional, and multimodal data. This challenge refers to methods for constructing meaningful and useful representations of state from such data.

Challenge 7*: High-level task planning*.

Robots will need to reliably execute high-level commands that fuse the previous six challenges to achieve a new level of utility, especially if they are to benefit the general public. For example, the command "get the milk" must autonomously generate the lower-level tasks of navigating to/from the refrigerator, opening/closing the door, identifying the proper container (milk containers may take many forms), and securely grasping the container. Loosely speaking, these challenges form a sort of "basis set" for the goals mentioned above. For example, human-like walking and running will rely heavily on 12 Challenges 1 (learning dynamics) and 2 (learning control policies), while teaching by demonstration will require advances in Challenges 4 (object recognition), 5 (interpreting human action), and 6 (sensor fusion).

Additional focus on applying Structures B, C, and D to robotics problems may very well catalyse significant advancement in many of the identified challenges.

Classifiers and discriminative models (Structure A) in robotics

The role of Structure A in robotics

Structure A involves using a deep learning model to approximate a function from sample input-output pairs. This may be the most general-purpose deep learning structure, since there are many different functions in robotics that researchers and practitioners may want to approximate from sample observations. Some examples include mapping from actions to corresponding changes in state, mapping from changes in state to the actions that would cause it, or mapping from forces to motions. Whereas in some cases physical equations for these functions may already be known, there are many other cases where the environment is just too complex for these equations to yield 14 acceptable accuracy. In such situations, learning to approximate the function from sample observations may yield significantly better accuracy. The functions that are approximated need not be continuous. Function approximating models also excel at classification tasks, such as determining what type of object lies before the robot, which grasping approach or general planning strategy is best suited for current conditions, or what is the state of a certain complex object with which the robot is interacting. The next section reviews some of the many applications for classifiers, regression models, and discriminative models that have appeared in the recent literature with robotics.

Generative and Unsupervised models (Structure B) in robotics

The role of Structure B in robotics

One of the characteristic capabilities that make humans so proficient at operating in the real world is their ability to understand what they perceive. A similar capability is offered in autoencoders, a type of deep learning model that both encodes observations into an internal representation, then decodes it back to the original observation. These models digest high-dimensional data and produce compact, low-dimensional internal representations that succinctly describe the meaning in the original observations . Thus, auto-encoders are used primarily in cases where high-dimensional observations are available, but the user wants a low-dimensional representation of state. Generative models are closely related. They utilize only the decoding portion of an autoencoder, and are useful for predicting observations. Inference methods may be used with generative models to estimate internal representations of state without requiring an encoder to be trained at all. In many ways, generative models may be considered to be the opposite of classifiers, or discriminative models, because they map from a succinct representation to a full high-dimensional set of values similar to those that might typically be observed.

Recurrent models (Structure C) in robotics

The role of Structure C in robotics

Recurrent models excel at learning to anticipate complex dynamics. The recurrent connections in such models give them a form of "memory" that they can use to remember the current state. This knowledge of state enables them to model the effects of time in a changing environment.

Policy learning models (Structure D) in robotics

The role of Structure D in robotics

Learning a near optimal (or at least a reasonably acceptable) control policy is often the primary objective in combining machine learning with robotics. The canonical model for using deep neural networks for learning a control policy is deep Q-learning . It uses a DNN to model a table of Q-values, which are trained to converge to a representation of the values for performing each possible action in any state. Although Structure D is quite similar to Structure A in terms of the model itself, they are trained in significantly different ways. Instead of minimizing prediction error against a training set of samples, deep Q-networks seek to maximize long-term reward. This is done through seeking a balance between exploration and exploitation that ultimately leads to an effective policy model. Ultimately, reinforcement learning models are useful for learning to operate dynamic systems from partial state information, and controllers based on deep reinforcement learning can be very computationally efficient at runtime [125]. They automatically infer priorities based on rewards that are obtained during training. In theory, they provide a complete control policy learning system, but they do suffer from extremely slow training times. Consequently, many of the works in the next section combine them with other approaches in order to seek greater levels of control accuracy and training speed

Conclusion

Deep learning has shown promise in significant sensing, cognition, and action problems, and even the potential to combine these normally separate functions into a single system. DNNs can operate on raw sensor data and deduce key features in that data without human assistance, potentially greatly reducing up-front engineering time. They are also adept at fusing high-dimensional, multimodal data. Improvement with experience has been demonstrated, facilitating adaptation in the dynamic, unstructured environments in which robots operate. Some remaining barriers to the adoption of deep learning in robotics include the necessity for large training data and long training times. Generating training data on physical systems can be relatively time consuming and expensive. One promising trend is crowdsourcing training data via cloud robotics . It is not even necessary that this data be from other robots, as shown by Yang's use of general-purpose cooking videos for object and grasp recognition . Regarding training time, local parallel processing and increases in raw processing speed have led to significant improvements. Distributed computing offers the potential to direct more computing resources to a given problem [88] but can be limited by communication speeds . There may also be algorithmic ways of making the training process more efficient yet to be discovered.

References:

[1].https://arxiv.org/ftp/arxiv/papers/1707/1707.07217.pdf

[2]. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436-444.

[3]. Jordan MI, Mitchell TM. Machine learning: arends, perspectives, and prospects. Science. 2015;349(6245):255-260.

[4] Böhmer W, Springenberg JT, Boedecker J, et al. Autonomous learning of state representations for control: an emerging field aims to autonomously

learn state representations for reinforcement learning agents from their real-world sensor observations. KI-Künstliche Intelligenz. 2015;29(4):353-362.

[5]. www.quora.com/How-does-deep-learning-work-for-robots.

[6]. www.quora.com/How-does-deep-learning-work-for-robots.

[7]. www.quora.com/How-is-deep-learning-used-in-robotics

[8].www.analyticsvidhya.com/blog/2018/05/nvidias-deep-learning-ai-trains-robots-to-copy-and-execute-human-actions

[9]. Deep Learning in Robotics: A Review of Recent Research Harry A. Pierson, Michael S. Gashler.