# RobotGPT: From ChatGPT to Robot Intelligence

Hongmei He*,
*School of Science, Engineering and Environment, University of Salford, Manchester, UK
Email: h.he5@salford.ac.uk

***Abstract*—ChatGPT, a language model developed by OpenAI, has triggered a new wave of AI development. This article reviews the principles of ChatGPT, analyzes five types of robot perception related to human perception (visual, auditory, tactile, olfactory, and gustatory), and seven types of robot intelligence (logical-mathematical, linguistic, spatial, musical, bodily-kinesthetic, interpersonal, and naturalistic), but excludes intrapersonal intelligence as a type of robot intelligence, distinguishing robots from human individuals. The author proposes a multimodal approach that takes advantage of GPT-3 to implement the seven types of robot intelligence, and provides a framework for implementing an intelligent robotic system, called RobotGPT.**

Keywords: ChatGPT, Generative Pre-trained Model, Tuning, Reinforcement Learning, Robot Perception, Robot Intelligence, RobotGPT.

## I. Introduction

Human: "What is ChatGPT?"
ChatGPT: "Hello! I'm ChatGPT, a language model developed by OpenAI. I was trained on a large corpus of text data using deep learning techniques, which enables me to understand natural language and generate responses to a wide range of questions and prompts... Let me know if there's anything specific you'd like to know or talk about!"

The above dialog was generated by a user and ChatGPT (Chat Generative Pre-trained Transformer), one of the most intelligent chatbots in the world. The main difference between ChatGPT and traditional human-computer interaction (HMI) lies in the way they process and respond to user input. Traditional HMI systems typically rely on pre-programmed rules and commands to interpret user input and generate limited responses. HMI often requires users to use certain keywords or adhere to a specific format in order to interact effectively with the system. In contrast, ChatGPT uses natural language processing (NLP) and machine learning (ML) techniques to understand and respond to user input in a conversational manner. ChatGPT can respond to a much wider range of inputs and questions than a conventional HMI.

ChatGPT has recently attracted the attention of the research community, the commercial sector, and the public. A wave of innovations based on ChatGPT has been triggered in various fields. For example, Gilson et al. [3] evaluated the performance of ChatGPT on questions in the United States Medical Licensing Examination Step 1 and Step 2 and analysed the answers for their interpretability by the user based on three qualitative metrics: logical justification of the chosen answer, presence of information within the question, and presence of information outside the question. The Microsoft Autonomous Systems and Robotics Group presented an experimental study on the use of ChatGPT for robotics applications [4].

As Zadeh stated in [6], humans have many remarkable capabilities, two of which are particularly important: (1) the capability to reason, converse and make rational decisions in an environment of imprecision, uncertainty, incompleteness of information, the partiality of truth, and possibility; (2) the capability to perform a variety of physical and mental tasks without measurement or calculation. There is ample evidence that machine intelligence has made some progress in the first capacity with advanced IoT and AI technologies. Although robotic cognitive capabilities are constantly evolving, we still have a long journey for the second capability of robots, which requires a machine brain capable of perceiving complex environments and making decisions promptly [7]. ChatGPT's advances in language capability, based on its extensively pre-trained generative language model and reinforcement learning capability, may provide a clue for the development of robot intelligence. ChatGPT integrated into the robot's brain could take robot intelligence a step forward for the two capabilities.

ChatGPT and Whisper APIs are now available [8]. With APIs, developers are allowed to integrate ChatGPT and Whisper models into their applications and products. Thus, it is possible to integrate ChatGPT into the brain of a robot to implement an intelligent robot system. This leverages ChatGPT's advances in flexible natural language and strong conversational capabilities and combines advanced perception and AI technologies.

This paper explores how ChatGPT capabilities can enhance robotic capabilities. Robotic systems, unlike text-only applications, require a deep understanding of real-world physics, environmental context, and the ability to perform physical actions. A generative robot model must have robust commonsense knowledge and a sophisticated world model, and be able to interact with the user and the surrounding environment to interpret and execute commands in a trustworthy way. These challenges go beyond the scope of language models, as they must not only understand the meaning of a given text but also translate intent into a logical sequence of physical actions. Therefore, we will explore the intersection of ChatGPT and an intelligent robot system and build a framework for robot intelligence taking the advances of ChatGPT.

## II. Principles of ChatGPT

Rapid advances in natural language processing (NLP) have led to the development of large-scale language models, such as BERT [9], GPT-3 [1], and Codex [10], which are innovating a variety of applications. These models have been successfully applied for various tasks such as text generation, machine translation, code synthesis, etc. Unlike previous models, which mostly use a single prompt, ChatGPT offers particularly impressive interaction capabilities through dialogues and combining text generation with code synthesis [4].

ChatGPT was developed based on the GPT-3.5 (Generative Pre-trained Transformer 3.5) architecture with three variants of 1.3B, 6B, and 175B parameters [5]. GPT-3.5 (i.e. InstructGPT) is trained on the same data sets as GPT-3, and fine-tuned using supervised learning and reinforcement learning with human feedback (RLHF) to incrementally improve the response to queries based on the extensively pre-trained generative language model (Figure 1). The multiturn dialogue interaction of ChatGPT supports a wide range of text- and code-based tasks.

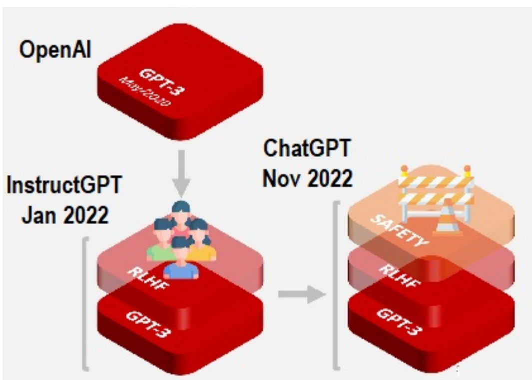


Fig. 1. From GPT-3 to ChatGPT [5]

### A. The Generative Pre-trained Transformer Model

GPT-3 has the same architecture as GPT-2, but with different conversions: the modified initialization, pre-normalization, and reversible tokenization, plus dense and locally banded sparse attention patterns, used alternately in the layers of the transformer [1]. To investigate the dependence of ML performance on model size, 8 different model sizes between 125M and 175B (the real GPT-3) were used. The smallest model (i.e., 125M) has 12 attention layers (Figure 2 (a)), each with 12 heads and 64 dimensions. The largest model has 96 attention layers with 96 heads and 128 dimensions. Table I summarises the parameters of the scaled architecture for training GPT-3. The batch size of GPT-3 increases and the learning rate decreases as the size of the GPT-3 model increases [11].

### B. Data Processing

The training data includes three types of data: the book data set, the web text data set, and the common crawl data set with 410 billion tokens. Data quality is one of the most important factors affecting the performance of an ML model. To improve data quality, the GPT-3 training data was processed in three steps (Figure 3).

(1) *feature extraction*: Data is automatically filtered based on similarity to a set of high-quality reference corpora. A tokenizer breaks a stream of characters into individual tokens and outputs a stream of tokens. HashingTF, a transformer, converts groups of terms into fixed-length feature vectors, linked to an index by applying a hash function.

(2) *Classification*: a logistic regression classifier is used to distinguish the curated data sets (WebText, Wikipedia, and the book corpus), representing positive examples, from the unfiltered common crawl data sets, representing negative examples. Moreover, a document from common crawl data sets is retained if it satisfies the condition: survival probability is greater than 1 - document score.

(3) *Deduplication*: To improve the data quality and prevent over-fitting, a fuzzy deduplication was performed at the document level to remove highly overlapping documents. Based on the document features generated by the classifier, a Spark Locality Sensitive Hashing was used with 10 buckets, so that similar documents end up in the same bucket.

The resulting common crawl data set and other high-quality data sets (WebText, book corpora, English Wikipedia documents) were mixed to form the final training data set. Samples in the test data that overlap with the training data are removed.

### C. Pre-training and Tuning

The training procedure consists of three phases [13]: (1) an unsupervised learning phase in which a powerful language model is trained on a large text corpus, (2) a fine-tuning phase to fit the model to a discriminative task with labeled data, (3) a reinforcement learning phase with human feedback (RLHF).

*Unsupervised Pre-Training* is to find a good initialization point of the model using unsupervised learning rather than to modify the supervised learning objective. Given a corpus of tokens $U = u_1, ..., u_n$, a standard objective is defined to maximize the likelihood:

$$L_1(U) = \sum_i logP(u_i|u_{i-k}, ..., u_{i-1}; \Theta), \tag{1}$$

where $k$ is the size of the context window and the conditional probability $P$ is modeled by a neural network with parameters $\Theta$, which are trained using stochastic gradient descent. The language model is a multi-layer transformer decoder, where, a multi-headed self-attention operation is applied to the input context tokens, followed by position-dependent feed-forward layers to produce an output distribution.

*Supervised fine-tuning* is used to adjust the parameters to a specific task. Consider a labeled data set $C$, consisting of a sequence of input tokens $x_1, ..., x_m$ along with a label $y$ as input to the pre-trained model to produce the activation value $h_l^m$ of the transformer, which is then fed into a linear layer with weights $W_y$ to estimate $y$:

$$P(y|x_1, ..., x_m) = softmax(h_l^m W_y). \tag{2}$$

thus to maximize the objective:

$$L_2(C) = \sum_{(x,y)} logP(y|x_1, ..., x_m). \tag{3}$$

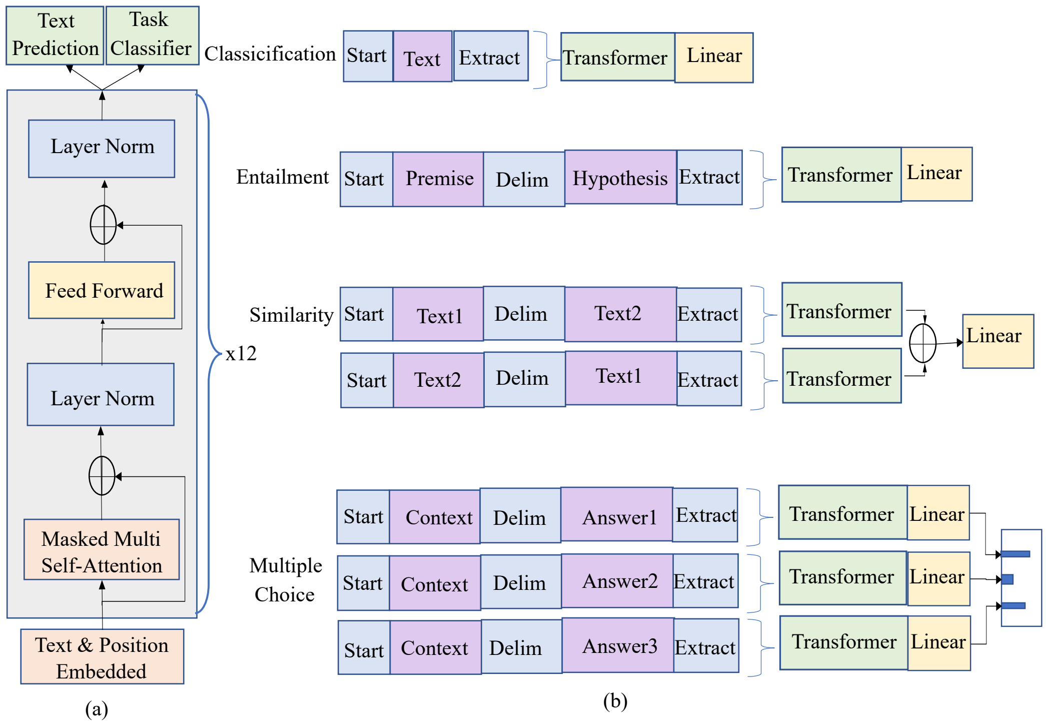


Fig. 2. GPT-3 model and tuning for various tasks.[13], [14]

TABLE I

SUMMARY OF SCALED ARCHITECTURE FOR TRAINING GPT-3. WHERE, $n_{paras}$ IS NUMBER OF PARAMETERS, $n_{layers}$ IS THE NUMBER OF LAYERS, $n_{models}$ IS THE NUMBER OF MODELS, $n_{heads}$ IS THE NUMBER OF HEADS, $d_{head}$ IS THE DIMENSIONS OF A HEAD [11]

| Model Names | $n_{paras}$ | $n_{layers}$ | $n_{models}$ | $n_{heads}$ | $d_{head}$ | Batch Size | Learning Rate |
|---|---|---|---|---|---|---|---|
| GPT-3 Small | 125M | 12 | 768 | 12 | 64 | 0.5M | $6.0 \times 10^{-4}$ |
| GPT-3 Medium | 350M | 24 | 1024 | 16 | 64 | 0.5M | $3.0 \times 10^{-4}$ |
| GPT-3 Large | 760M | 24 | 1536 | 16 | 96 | 0.5M | $2.5 \times 10^{-4}$ |
| GPT-3 XL | 1.3B | 24 | 2048 | 24 | 128 | 1M | $2.0 \times 10^{-4}$ |
| GPT-3 2.7B | 2.7B | 32 | 2560 | 32 | 80 | 1M | $1.6 \times 10^{-4}$ |
| GPT-3 6.7B | 6.7B | 32 | 4096 | 32 | 128 | 2M | $1.2 \times 10^{-4}$ |
| GPT-3 13B | 13.0B | 40 | 5140 | 40 | 128 | 2M | $1.0 \times 10^{-4}$ |
| GPT-3 175B (GPT-3) | 175.0B | 96 | 12288 | 96 | 128 | 3.2M | $0.6 \times 10^{-4}$ |

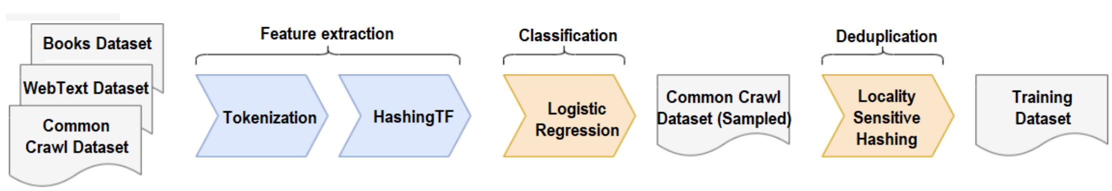


Fig. 3. Flowchat of Data Processing[11]

Figure 2 (b) shows fine-tuning transformations for four tasks. In an input sequence, two tokens are connected by a delimiter token ($). For classification, all structured inputs are transformed into token sequences, proceeded by the pre-trained model, followed by a Linear + Softmax layer. For entailment, the token sequences are a concatenation of premise $p$ and hypothesis $h$. For similarity, the two propositions are compared, with arbitrary order of input sequences, two produced representations $h_l^m$ are added to feed the linear output layer. For multiple-choice, outputs from multiple linear layers are selected in terms of their scores.

*Reinforcement learning with human feedback* (RLHF) The GPT text model created in the two steps above is further refined using reinforcement learning with human feedback to achieve better results. During the training process, the human trainer acts as both user and an AI assistant. In this phase, a context document $z$, a user-posed question $q$, and a set of possible answers $a_k$ are combined to $\{z, q, \$, a_k\}$. Each of these sequences is processed independently with the model and then normalized by a Softmax layer, thus producing an output distribution over the possible answers to the user's question.

Ouyang et al. [15] divide the tuning method into three steps: supervised fine-tuning, reward model training, and reinforcement learning. Figure 4 illustrates the tuning method, in which the reinforcement learning phase is implemented by proximal policy optimization (PPO) on a reward model, trained in step 2, boxes A-D are samples from the models, ranked by labels, and blue arrows indicate that the corresponding data is used to train the target models.

## III. Robot Perception and Intelligence

### A. Robot Perception

Human perception is a complex process involving sensory processing and cognitive interpretation. Our perceptual experiences are influenced by a variety of factors, including our past experiences, attention, motivation, and expectations. As a result, different people may perceive the same sensory information in different ways, and our perceptions may be subject to biases and distortions. Similarly, robot perception refers to a robot's ability to sense and interpret its environment using sensors and artificial intelligence (AI) algorithms. Perception is a critical capability for robots because it enables them to interact with the world in meaningful ways and perform tasks autonomously. We can define robot perception as follows:

$$Robot perception = sensing + interpreting = f(S), \quad (4)$$

where $S$ is the readings of a sensor array (homogeneous or heterogeneous sensors). Like human perception, robots could also have five essential types of perception.

*Visual perception* (VP) refers to a robot's ability to perceive and interpret visual information from cameras and other visual sensors. Key processes include image processing, object recognition, scene understanding, visual tracking, 3D perception.

*Auditory perception* (AP) refers to a robot's ability to perceive and interpret sounds and speeches. *AP* enables robots to understand and respond to verbal commands and other auditory cues. Key processes include speech recognition, sound localization, acoustic scene analysis, speech synthesis, and auditory feedback.

*Tactile perception* (TP) refers to the robot's ability to perceive and interpret tactile information through touch, using tactile sensors attached to the robot's fingers, hands, and other body parts, for detecting pressure, force, temperature, and other physical properties, which are then used to create a model of the objects with which the robot interacts.

*Olfactory perception* (OP) refers to a robot's ability to perceive and interpret odors in the environment. This is achieved through the use of electronic noses, which are arrays of gas or chemical sensors that can detect and analyze the chemical composition of odors.

*Gustatory perception* (GP) refers to a robot's ability to simulate the human sense of taste by detecting and distinguishing different tastes in foods or beverages. It uses electronic sensors and algorithms to replicate the process of human taste perception. Currently, *GP* is in the early research stages, and accuracy and sensitivity can still be improved.

### B. Robot Intelligence

Different areas of the human brain correspond to specific knowledge spaces, all of which are distinct and relatively independent of each other. These knowledge spaces originate in human cognition and give rise to human intelligence, a multifaceted concept that encompasses a wide range of cognitive abilities and skills. Most intelligence theories divide human intelligence into eight categories: linguistic, logical-mathematical, spatial, physical-kinesthetic, musical, interpersonal, intrapersonal, and naturalistic intelligence [16]. Accordingly, we define robot intelligence (*RI*) in terms of human intelligence but excluding intrapersonal intelligence. Human intrapersonal intelligence is often associated with self-awareness, personality, individual preferences, bias, and ethics. However, robots as tools to augment human cognition are not expected to have intrapersonal intelligence. Self-awareness could be useful, but individual biases and ethical issues should be consciously considered.

*Linguistic Intelligence* (*LI*) refers to a robot's ability to understand and process human language, both written and spoken, in such a way that it can communicate effectively with humans. To achieve linguistic intelligence, robots must be equipped with natural language processing (NLP) capabilities. This involves using AI algorithms to analyse and understand human language, including its syntax, grammar, and semantics, interpret spoken or written language, and generate appropriate responses. Therefore, the robot's linguistic intelligence could be denoted as $LI = f(S|speechortext)$, where $S$ represents speech signals from acoustic sensors or text by input or from a camera.

*Logical-mathematical intelligence* (*LmI*) refers to a robot's ability to reason logically, solve problems, and perform mathematical operations using algorithms and computation. To achieve logical-mathematical intelligence, robots must be equipped with sensors, processors, and algorithms that enable robots to collect, process and analyse data, and programmed to follow logical rules and perform calculations in a systematic and efficient manner. It could be denoted as $LmI = f(S|K)$, $S$ is the readings from the sensors and $K$ is the knowledge base. This means that logical-mathematical intelligence is the function of the sensor readings $S$ under the condition of knowledge $K$.

*Spatial intelligence* (*SI*) refers to a robot's ability to understand and navigate three-dimensional space using sensors and AI algorithms. To achieve spatial intelligence, robots must be equipped with a variety of sensors, such as cameras (*C*), lidar (*L*), sonar (*S*), radar (*R*), infrared (*IR*), etc. Advanced sensor technologies enable robots to perceive their environment, and AI algorithms enable them to analyse and interpret sensor data, create internal maps of their environment, and navigate through complex environments. Therefore, $SI = f(C,L,S,R,IR|K)$.

*Bodily-kinesthetic intelligence* (*BkI*) refers to a robot's ability to control its movements and manipulate objects in a coordinated and precise manner. To achieve bodily-kinesthetic

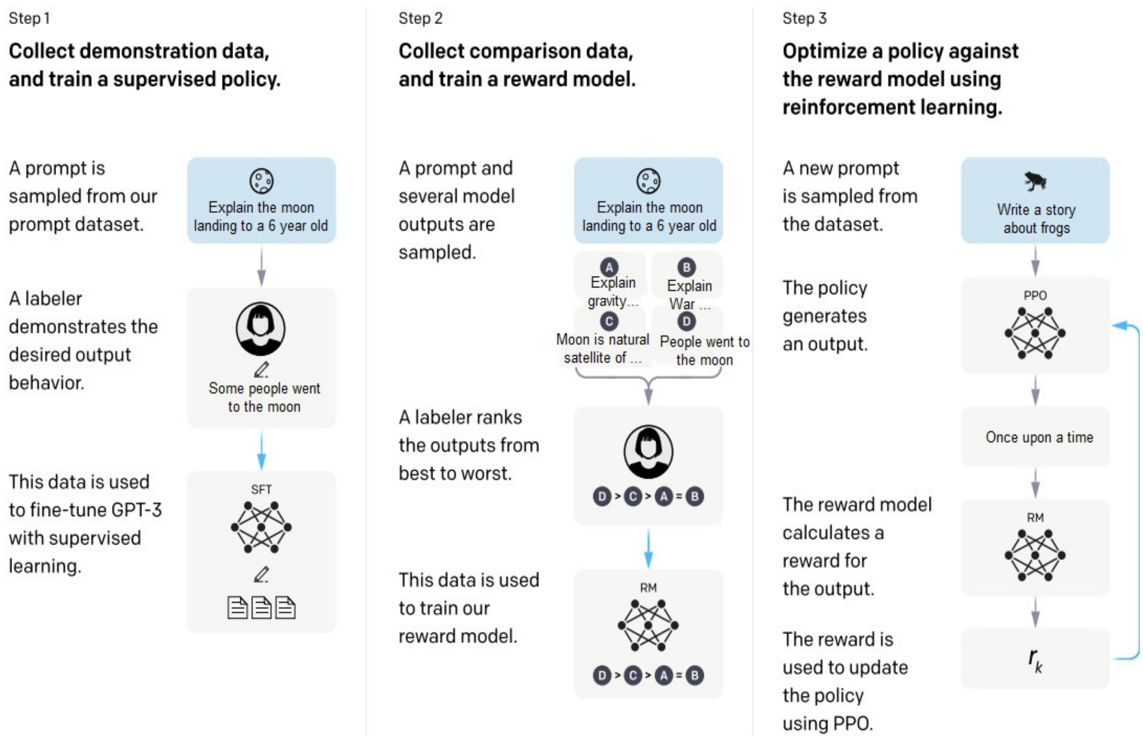


Fig. 4. Three Steps of Tuning [15]

intelligence, robots must be equipped with a variety of sensors, such as force sensors ($F$), tactile sensors ($T$), and vision sensors ($V$), etc. which allow them to sense their environment and the objects within it. Robots must also be equipped with actuators, such as motors, which allow them to move their limbs and manipulate objects precisely, Therefore, $BkI = f(F,T,V|K)$.

*Musical intelligence* (*MI*) refers to a robot's ability to create, understand, and perform music. To achieve musical intelligence, robots must be equipped with a variety of sensors, such as microphones ($M$) and accelerometers ($A$), which enable them to perceive sound and motion. Embedded AI algorithms enable them to analyse and interpret musical elements such as rhythm, melody, and harmony. Therefore, $MI = f(M,A|K)$.

*Interpersonal intelligence* (*IeI*) refers to a robot's ability to interact with humans in a way that is socially and emotionally intelligent. To achieve interpersonal intelligence, robots must be equipped with natural language processing capabilities and AI algorithms that enable them to interpret and respond appropriately to human emotions. They must also be equipped with physical features and movements that are comprehensible and non-threatening to humans. Therefore, $IeI = LI + EI + SI$, where *LI* is linguistic intelligence, *EI* is emotional intelligence, and *SI* is spatial intelligence.

*Naturalistic intelligence* (*NI*) refers to a robot's ability to understand and interact effectively with the natural world, including skills such as observation, classification, and identification of living and nonliving things. To achieve naturalistic intelligence, robots must be equipped with a variety of sensors, such as cameras ($C$), microphones ($M$), and environmental sensors ($E$) (e.g., temperature, humidity, CO2, etc.), which enable them to perceive the natural world. AI algorithms enable them to analyse and interpret natural patterns and processes and respond accordingly. Therefore, $NI = f(C,M,E|K)$.

### C. Generative Pre-Trained Model

Robots and autonomous systems (RAS), equipped with AI technologies, seek to mimic the adaptive and intelligent capabilities of human problem solving [12], improve the ability to perceive complex environments, understand human behaviours and/or questions, and make decisions quickly by stimulating their cognitive and developmental abilities in a trustworthy manner toward human intelligence [7].

To achieve robot intelligence, they must be equipped with sensors, processors, and algorithms that enable them to collect, process, and analyse data. Seven generative pre-trained models need to be produced with respect to the seven types of robot intelligence.

$$RI = F(f(S), g(K)),\ or\ RI = Z(f(S)|g(K)). \tag{5}$$

where $S$ is a set of samples composed of readings from sensor arrays for different perceptions, $K$ is a set of knowledge. This means that a type of robot intelligence is either the fusion of perceptions from the sensor array and the knowledge base or the function of the perception of the sensor array under the condition of knowledge $K$. Therefore, seven generative pre-trained models with respect to the seven types of robot intelligence could be obtained using the method of GPT-3, but the training data includes a knowledge base and robot perception, required by the corresponding robot intelligence, described in Section III-B.

After training, to tune the seven generative pre-trained models, seven tuning processes are performed with respect to the specific tasks for different types of intelligence. Seven award models are trained for reinforcement learning, like GPT-3, but with the perception information in the surrounding environment and the feedback from humans/robots in the team.

### D. Reinforcement learning of robots

Reinforcement learning is based on classical reinforcement learning with a reward model incorporating perceptions of the sensor array on it and the feedback of the human-robot team that the robot is working with, whereas ChatGPT is on human feedback only. In addition, the policy is constrained by legal compliance with respect to corresponding standards, regulations, law enforcement, and ethics. Currently, ChatGPT has drawn much attention to legal compliance.

### E. RobotGPT

With respect to the above analysis, seven generative pre-trained models are required to implement seven types of robot intelligence. ChatGPT provides a good methodology for implementing robot intelligence, although it is more complicated than GPT-3, which only allows text conversation. GPT-4, released on 14 Mar 23, is more powerful than GPT-3 because it provides a more accurate solution to a query and supports image understanding [17]. The approach of GPT-3 or GPT-4 can be extended to develop the seven types of robot intelligence. A framework for robot intelligence is proposed as shown in Figure 5, denoted as RobotGPT.

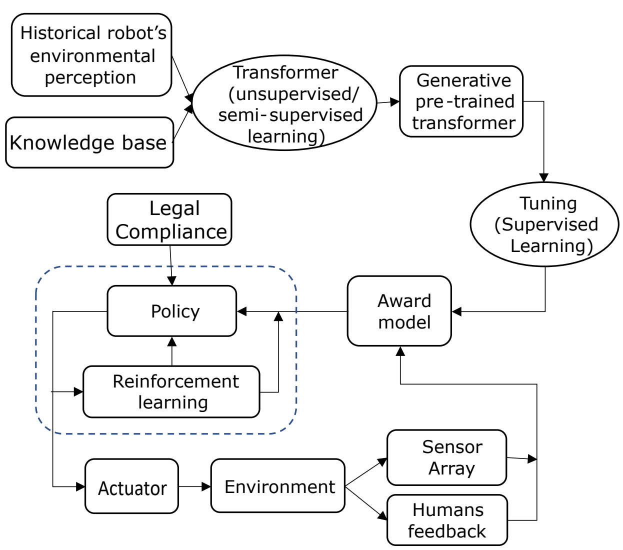


Fig. 5. Framework of RobotGPT [15]

## IV. Conclusions

This article provides an overview of the principle of GPT-3. ChatGPT provides a methodology that can be extended for the development of robot intelligence. In terms of human perception and human intelligence, the author divides robot perception into five categories and robot intelligence into seven types of robot intelligence, excluding intrapersonal intelligence, which distinguishes human individuals from robots. Based on this analysis, the author provides a multimodal approach and uses the approach of GPT-3 to implement the seven types of robot intelligence, and provides a framework for implementing robot intelligence (called RobotGPT). This work needs further analysis and development. It is an extensive technical project. The main challenge is to recognise the real world, which is complicated, dynamic, and uncertain.

## References

[1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. In: Advances in neural information processing systems, 33 (2020), pp. 1877-1901.

[2] A. Farseev. Is Bigger Better? Why The ChatGPT Vs. GPT-3 Vs. GPT-4 'Battle' Is Just A Family Chat. 17 Feb. 2023. www.forbes.com. Accessed on 11/03/2023.

[3] Gilson A, Safranek CW, Huang T, Socrates V, Chi L, Taylor RA, Chartash D How Does ChatGPT Perform on the United States Medical Licensing Examination? The Implications of Large Language Models for Medical Education and Knowledge Assessment. JMIR Med Educ 2023;9:e45312.

[4] S. Vemprala, R. Bonatti, A. Bucker and A. Kapoor. ChatGPT for Robotics: Design Principles and Model Abilities. 20 Feb 2023. https://www.microsoft.com/en-us/research/group/autonomous-systems-group-robotics/articles/chatgpt-for-robotics/. Accessed on 10/03/2023.

[5] A. Mandour, GPT-3.5 model architecture, https://iq.opengenus.org/gpt-3-5-model/. Accessed on 17/03/2023.

[6] L. A. Zadeh, "Toward Human Level Machine Intelligence - Is It Achievable? The Need for a Paradigm Shift," in IEEE Computational Intelligence Magazine, 3(3), pp. 11-22, August 2008.

[7] H. He, J. Gray, A. Cangelosi, Q. Meng, T. M. McGinnity, J. Mehnen, The Challenges and Opportunities of Human-Centred Artificial Intelligence for Trustworthy Robots and Autonomous Systems, IEEE Transactions on Cognitive and Developmental Systems, 14(4), 2022, pp.1398-141202.

[8] G. Brockman, A. Eleti, E. Georges, J. Jang, L. Kilpatrick, R. Lim, L. Miller and M. Pokrass.Introducing ChatGPT and Whisper APIs. https://openai.com/blog/introducing-chatgpt-and-whisper-apis, accessed on 10/03/2023.

[9] J. Devlin, M-W Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv:1810.04805, 2018.

[10] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al. Evaluating large language models trained on code. arXiv:2107.03374, 2021.

[11] GPT-3 An Overview, dzlab. https://dzlab.github.io/ml/2020/07/25/gpt3-overview/. Accessed on 11/03/2023.

[12] E. E. Alves, D. Bhatt, B. Hall, K. Driscoll, A. Murugesan, and J. Rushby. Considerations in assuring the safety of increasingly autonomous systems. NASA/CR-2018-220080, NASA, July 2018.

[13] Alec Radford, K. Narasimhan, T. Salimans and I. Sutskever. Improving Language Understanding by Generative Pre-Training, 2018. https://paperswithcode.com/paper/improving-language-understanding-by. Accessed on 12/03/2023.

[14] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei and I. Sutskever. Language models are unsupervised multitask learners, OpenAI blog 1.8 (2019): 9.

[15] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askellâ€ , P. Welinder, P. Christiano, J. Leike, and R. Lowe. Training language models to follow instructions with human feedback. arXiv:2203.02155v1 [cs.CL] 4 Mar 2022.

[16] A talent for every child: The eight types of intelligence, https://www.iberostar.com/en/inspiration-guide/wellness/eight-types-of-intelligence/. Accessed on 12/03/2023.

[17] https://openai.com/research/gpt-4. accessed on 17 Mar 2023.