



The five Is: Key principles for interpretable and safe conversational AI

Downloaded from: <https://research.chalmers.se>, 2026-04-14 00:29 UTC

Citation for the original published paper (version of record):

Wahde, M., Virgolin, M. (2021). The five Is: Key principles for interpretable and safe conversational AI. ACM International Conference Proceeding Series: 50-54.
<http://dx.doi.org/10.1145/3507623.3507632>

N.B. When citing this work, cite the original published paper.



The five Is: Key principles for interpretable and safe conversational AI

Mattias Wahde*

Marco Virgolin

mattias.wahde@chalmers.se

marco.virgolin@chalmers.se

Chalmers University of Technology

Gothenburg, Sweden

ABSTRACT

In this position paper, we present five key principles, namely *interpretability*, *inherent capability to explain*, *independent data*, *interactive learning*, and *inquisitiveness*, for the development of conversational AI that, unlike the currently popular black box approaches, is transparent and accountable. At present, there is a growing concern with the use of black box statistical language models: While displaying impressive average performance, such systems are also prone to occasional spectacular failures, for which there is no clear remedy. In an effort to initiate a discussion on possible alternatives, we outline and exemplify how our five principles enable the development of conversational AI systems that are transparent and thus safer for use. We also present some of the challenges inherent in the implementation of those principles.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; **Natural language interfaces**; **Collaborative interaction**.

KEYWORDS

Interpretable AI, conversational agents, dialogue managers, multimodal human-computer interaction

ACM Reference Format:

Mattias Wahde and Marco Virgolin. 2021. The five Is: Key principles for interpretable and safe conversational AI. In *2021 The 4th International Conference on Computational Intelligence and Intelligent Systems (CIIS 2021)*, November 20–22, 2021, Tokyo, Japan. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3507623.3507632>

1 INTRODUCTION

The field of conversational AI is in rapid development and it has important applications in healthcare (for example elderly care), education, business (for instance customer service), and so on [2,

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIIS 2021, November 20–22, 2021, Tokyo, Japan

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8593-0/21/11...\$15.00

<https://doi.org/10.1145/3507623.3507632>

15, 25]. Those applications involve the use of conversational agents, which are systems intended for natural, multimodal interaction with human users, using text, speech, touch, gestures, and so on.

At present, research in conversational AI is strongly dominated by black box approaches, in which deep neural networks (DNNs) are used to encode statistical language models that underlie the dialogue capabilities of a conversational agent. DNNs have achieved spectacular success in many AI subfields, notably image processing and speech processing [10]. Thus, extending their use to conversational AI may seem natural, and many promising results have indeed been presented also in the case of conversational AI [19].

However, there are considerable drawbacks with using black box DNN-based systems, especially in high-stakes applications (such as, for example, healthcare) [22, 25]. Due to their distributed nature, black box systems must generally be trained with machine learning methods (deep learning) that require *massive amounts of data*. Because the data sets are so large, it is nearly impossible to carry out a quality check to discard, e.g., unintended biases and sensitive content [3, 4]. Therefore, for conversational AI, the use of DNNs can result in conversational agents that, in addition to being opaque, are unsafe for use, something that is discussed in more detail in 2 below.

While evidence on the unsuitability of data-hungry methods for safe conversational AI is mounting, a large part of current research focuses on creating methods to explain why black boxes make certain decisions, in an attempt to remedy the lack of transparency. However, explanation methods have their own limitations [18, 22].

To the best of our knowledge, no proposals exist yet on what could be radically different ways to move forward to create conversational AI that is transparent and accountable by construction. In this position paper, we initiate this discussion by proposing *five key design principles* that next-generation conversational AI should adhere to, in our view, to ensure transparency and accountability. These are *interpretability*, *independent data*, *inherent capability to explain*, *interactive learning*, and *inquisitiveness*¹; We call these the “five Is” of next-generation conversational AI; see Fig. 1. In short, *interpretability* and *inherent capability to explain* are necessary aspects to ensure transparency and accountability while the conversational agent is being built and during use, respectively. *Independent data* means that the agent should be easily adaptable to different application-specific knowledge bases, without the need of entirely re-working the way it handles information processing. Last but not least, *interactive learning* and *inquisitiveness* represent

¹Here, we use the word *inquisitiveness* in its positive sense, conveying a meaning similar to the curiosity of a child.

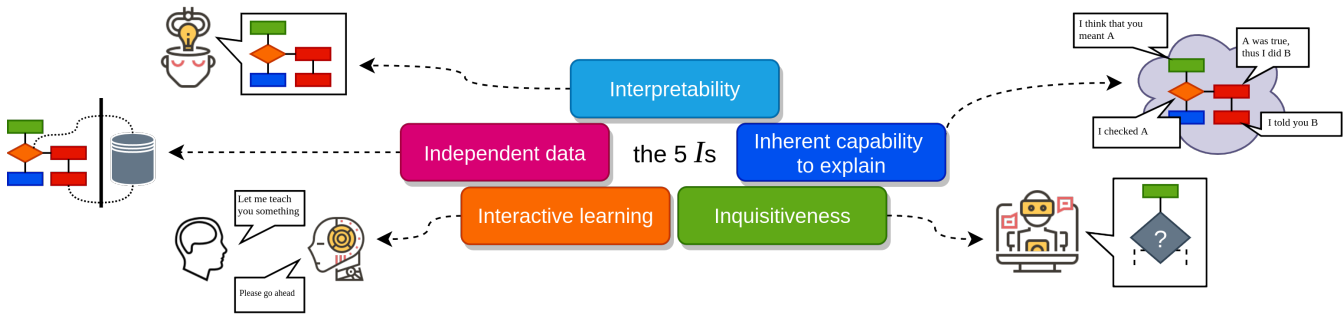


Figure 1: Five key principles for interpretable and safe conversational AI.

our proposal to expand the arsenal of available training methods by incorporating a modality of learning where the need for collecting and curating large amounts of data is replaced by a process that is natural to humans.

The paper can be read as a kind of manifesto, emphasizing the five design principles listed above.

2 BACKGROUND & RELATED WORK

Black box systems such as DNNs have shown exceptional capability to capture and reproduce the statistical properties of language, at times generating documents and dialogues that are indistinguishable from those produced by humans [1, 5]. However, there are several reasons for expressing concern with the use of black box systems such as DNNs in conversational AI.

First of all, with DNNs, there is a near-complete opaqueness of the underlying decision-making required for the conversational agent to generate an utterance. In fact, the decision-making of a DNN is the result of the concerted actions of myriads of neurons. For example, the state-of-the-art DNN-based language model GPT-3 uses 175 billion parameters [5]. Even if such a system can generate sensible answers to many queries, it is far from guaranteed to do so reliably [17], and it is nearly impossible to foresee occasional spectacular failures [9]. Furthermore, explanations of the underlying decision-making in such systems are generally not predictive in nature, but must rather be carried out on a case-by-case basis, *ex post facto* [18, 22].

Second, as briefly mentioned above, black box dialogue systems generally require huge amounts of data for their training [3]. In the case of so called *chatbots* (conversational agents intended for casual discussions on everyday topics) one can source almost unlimited amounts of dialogue data from the internet, e.g., twitter feeds. More often that not, however, such data are of rather dubious quality and may contain significant biases, involving language that is violent, racist, sexually explicit and so on [3]. One of the supposed advantages of DNNs, being end-to-end systems, is their ability to learn directly from data, without requiring any human involvement (e.g., hand-coding) beyond preparing the data, the system architecture, and the training algorithm. Yet that argument misses the extremely important point that the data sets, which are generally huge, must be carefully curated when using black box systems, as it is impossible to check whether they inadvertently learn *unwanted* biases during training [3]. It can thus be argued that the advantages of

using such systems evaporate to a large extent, considering the massive amounts of work required to check the data.

A clear example of this is the case of Tay [21], a chatbot released on the internet by Microsoft in 2016, which was able to learn from its interaction with thousands of users. Because of a lack of quality control in the interactions, within a few hours Tay developed nasty traits, expressing racist and sexist views, forcing permanent shut-down of the agent. Furthermore, it was recently shown that since DNNs can memorize part of the training data within their parameters, it is possible to carry out cyber-attacks to retrieve potentially privacy-sensitive information initially present in the data [6].

An even more serious issue appears when considering the use of black box systems for so called *task-oriented agents* that, unlike chatbots, are intended to provide precise, consistent, and detailed interaction with users regarding specific topics (e.g., in healthcare [16]), thus requiring significant knowledge on those topics, beyond the mere general knowledge expected from a chatbot. A recent disastrous example concerns the use of GPT-3 as a conversational agent intended to provide assistance for people with mental health problems. Here, a user (thankfully only *simulating* a patient suffering from depression) was *advised* by the black box system to commit suicide [9]. Thus, it is clear that black box systems cannot just be naïvely adopted and expected to make no mistakes. In all fairness, it is often the case that the authors of black box models warn about these problems (see, e.g., [5]).

Third, the prevailing (and rising) legal trends do *not* favor black box systems: There is proposed legislation both in the US (the algorithmic accountability act) and in the EU (the right to an explanation) that clearly favors more transparent and accountable approaches than those offered by black box systems. Incidentally, this is also a case where most of the ongoing research is at odds with the desires and requirements of end users, for whom transparency and accountability are clearly favored [12].

One may thus wonder whether other research avenues exist, which are better suited for transparent and accountable conversational AI. In fact, preceding the advent of black box systems, many different approaches to dialogue systems were considered, ranging from simple template-matching and finite-state systems [13, 26], to much more sophisticated (but also less transparent) approaches such as those based on partially observable Markov decision processes (POMDPs) [28]. At the time, transparency and accountability

were not as relevant and were not explicitly addressed. This essentially translates to the fact that, today, simpler approaches enjoy better transparency mostly because they are in great part hand-crafted and remain somewhat limited, and not because they were explicitly designed to facilitate transparency and accountability. Other modern approaches exist regarding language processing that are very different from DNN-based ones (see, e.g., [8, 23, 24]), however these do not have transparency as a primary focus either. We argue that one should strive to design systems that are *both* transparent *and* complex enough to enable human-like interactions. To that end, we believe that it is of primary importance to study what should be the design principles to ensure that transparency will be preserved as a conversational system grows increasingly complex.

3 DESIGN PRINCIPLES FOR TRANSPARENT CONVERSATIONAL AI

We propose five key design principles that, in our view, should be incorporated in conversational systems in order to ensure transparency, accountability, and safe use.

3.1 Interpretability

Contrary to the principles of black box systems, we propose, as the first design principle, the use of interpretable primitives in conversational systems, i.e. components that perform high-level operations such that both their purpose and mode of operation are easily human-understandable. Examples of such components are program constructs such as sorting operations, fetch commands, comparison conditions, and so on, which can be combined to carry out very complex operations that still remain human-readable.

We also suggest following the so called *pipeline model* (see e.g. [25]) whereby input processing, cognitive processing (decision-making), and output processing are clearly separated steps. For the case of input processing, various options could be used, even black box models, as long as they are ultimately mapped to a set of explicit intention sentences, which act as entry points for the cognitive processing. In cases where the input does not perfectly match any such intention sentence (which can happen, for example, with an approximation provided by a black box model), the agent should notify the user that it may require a clarification.

The output processing (that follows the cognitive processing) should also involve clearly defined, interpretable output sentences but, in this case, a semantic grammar or a black box model could be used for *altering* the exact formulation of the output, without changing its semantics, thus allowing a more life-like interaction.

The crucial middle step (cognitive processing), where the agent carries out deliberation and decision-making should, in our view, be implemented in a generic fashion (exemplified below), using a set of interpretable primitives that make it easy to follow (and correct, when necessary) every step of the agent’s decision-making as well as allowing re-use of implemented sequences of such primitives in other agents. A simple example may help to illustrate this process. Consider an agent that is capable to answer questions on geography and demographics, thus containing a database with information about countries, cities, and so on. Now, in this example, the user asks the agent *which is the second largest city in France?* After identifying the user’s input (details omitted here), the agent then carries

```
out (in pseudo-code) the following generic cognitive processing
FindAll(country=France, category=city)→List1,
SortAscending(List1,population)→List2,
GetElement(List2, 2)→Element,
ExtractValue(Element, name)→cityName.
```

In plain English, the agent thus starts by finding all data items pertaining to cities in France, then sorts the resulting list in ascending order based on population, then extracts the second element from the list and, finally extracts the name property from that element, storing it such that it can be accessed by the output processing step that follows the cognitive step. Note that the sequence of primitives, in addition to being transparent, is completely generic: The same sequence (but with other input parameters) could be used, say, in an on-board information system on a bus, for answering questions of the form *What is the second stop after this?*

3.2 Inherent capability to explain

Another key principle is that a conversational agent should be able to *explain* its reasoning, in a non-technical manner, allowing both developers and end users to understand how the agent reached a particular conclusion. This ability is particularly important in cases where the agent’s conclusion differs from what the user expected. As an example of such a situation, one can consider an agent that offers decision support to a driver (of a truck, forklift, or any other vehicle). In a case where the agent suggests a counter-intuitive maneuver, obtaining a clear explanation may be crucial.

With the approach described in Subsection 3.1 above, where the cognitive processing is generated as a sequence of elementary, generic primitives, it is quite straightforward to include the ability to explain, since every primitive of this kind can be associated with a description of itself, taking into account also the input parameters. Continuing on the example given in Subsection 3.1, where the agent answers the question *what is the second largest city in France?*, the automatically generated explanation (from the agent) could read: *First, I found all data items pertaining to cities in France. Then, I sorted them in ascending order based on population. Next, I got the second element from that list, and extracted its name property.*

3.3 Independent data

We further propose that an agent’s memory should be divided into two clearly separated parts: The *procedural memory* that encodes the dialogue capabilities described in the previous subsections, and the *declarative memory* that contains the facts known to the agent.

Along with the generic cognitive processing described above, this clear separation makes it possible to replace parts (or all) of the declarative memory, e.g. application-specific parts, without having to change the procedural parts. As an example, for an agent used as a museum guide, one can easily just replace the declarative memory in order to use the agent at a different museum, or even in a different task, e.g., as a city guide or a tour guide. Thus, with this approach, for many applications it is possible to use an agent in an off-the-shelf manner, just providing it with a suitable declarative memory.

3.4 Interactive learning

In the early days of conversational AI, agents were typically generated using hand-coding, i.e. a process whereby the developer would implement, by hand, the dialogue capabilities of the agent. With the advent of black box systems, this process was largely replaced by an automated approach involving machine learning using large dialogue data sets for training. The supposed advantage with the latter approach is that, in avoiding hand-coding, one can generate better performing conversational agents, without time-consuming manual implementation. However, as noted in Section 2, the machine learning approach has its own drawbacks, not least the need for careful checking of (vast amounts of) data, to avoid unwanted biases.

Here, as the fourth key design principle, we argue that interactive learning whereby a person (acting as a tutor) teaches a machine new capabilities [20, 27], should be a primary approach when building conversational agents. Coupled with the inquisitiveness (described below), this method has the advantage of resembling the manner in which humans, not least children, learn new concepts. Just like hand-coding, this approach also gives the developer full control over what the agent learns but, unlike hand-coding, interactive learning makes it possible also for a non-expert to teach an agent new capabilities. Moreover, similarly to what was described in Sect. 3.2, the interpretable primitives could be equipped with simple sentences by which they can be identified in a learning situation. A specific example (cf. the example in Sect. 3.1) should serve to illustrate the principle:

User: *I will teach you how to answer the question “what is the n^{th} largest city in a country x ”.*

Agent: *OK, please go ahead.*

User: *First, find all city data items in x .*

Agent: *Yes.*

User: *Then sort the items in ascending order based on population.*

Agent: *Understood.*

User: *Next, find the n^{th} element.*

Agent: *OK.*

User: *Then extract and return the name property from that element.*

Agent: *All right.*

At the end of this process, the agent will have assembled part of the cognitive processing as per the example of Sect. 3.1.

3.5 Inquisitiveness

The final key design principle involves an agent’s ability to actively seek information and thus to expand its capabilities, for example by attempting to make connections between different items in the data, categorizing data, and so on, in interaction with the user. Here, the agent should make *hypotheses*, but not commit anything to its memory without first checking with the user, thus making sure that the agent will not accidentally learn things that it should not. The inquisitiveness will also reduce the risk of accidentally omitting important aspects during the interactive learning: The agent should act as an active, eager student, rather than just passively receiving information. A simple example could be as follows:

Agent: *I know that the planet Mars has two moons. I also know of another planet called Venus. Does it have moons?*

User: *No it does not.*

Agent: *OK. What about the Earth?*

User: *It has one moon.*

Agent: *Understood.*

4 DISCUSSION

We have presented five design principles that we believe are key for the development of interpretable and safe conversational AI. While we hope that many readers will agree with the proposed principles, and some may even find them very natural, we remark that the currently popular black box models are ill-suited to adhere to these principles.

An important objection that might be raised against the position taken in this paper is the fact that black box approaches display remarkable performance, partly due to the fact that such approaches are eminently suited for current hardware, e.g., GPUs [11], while it is far from clear whether systems based on the principles proposed here would be able to reach similar performance, given current computers. However, as argued in [22], there is no fundamental reason to believe that one would have to sacrifice performance to achieve better interpretability; at the very least, we argue that this issue should be thoroughly investigated.

Interpretability is meant to make the inner workings of the system as explicit and transparent as possible to those involved in the design of the agent. By exposing interpretable primitives and their connections that together define the cognitive processing, an agent built in this manner makes it possible for the developers to get a clear overview of the system and thus to track, understand, and, if needed, correct its actions. By contrast, the inherent capability to explain is meant to provide virtually *any* user interacting with an agent with a clear explanation of the agent’s reasoning. The set of primitives need not be sufficiently expressive to carry out all forms of computation, let alone achieve general AI; the primitives need only enable the conversational agent to carry out its task, while remaining sufficiently high-level so that even a combination of rather a large number of primitives would remain interpretable. The exact nature of the primitives (as well as the number of primitives required) is an important issue for further work and will most likely depend on the application at hand.

The independent data principle encourages the design of conversational agents that are as general and decoupled from any particular application as possible, thus also making it possible to re-use existing cognitive capabilities across different application domains. Interactive learning and inquisitiveness are intended to provide a natural process for defining and improving conversational agents. Systems built according to these principles provide an interface for defining, expanding, and correcting both the cognitive capabilities of the agent and its application-specific knowledge, making it possible for non-experts to seamlessly train a conversational agent.

There exists a number of challenges for the development of conversational AI systems that adhere fully to these five principles. For example, it should be investigated what set of high-level primitives would be required for the agent to be taught complex procedures via interactive learning. Investing research efforts into interactive learning methods and inquisitiveness routines is likely to be of crucial importance for developing systems that adhere to the principles introduced here. Such methods need not themselves

be interpretable (see, e.g., [20]), as long as their final output, i.e. the conversational agent, is interpretable.

Living in an age of big data, however, it would also make sense to try to automatically infer some of the basic cognitive capabilities, in the form of combinations of the interpretable primitives, directly from data. We remark that, as long as interpretability is maintained, it will be possible (and necessary) to check and correct for the pick-up of unwanted biases when learning automatically from data prior to using the system in practice. Note that, as soon as non-differentiable primitives are used (e.g., IF-THEN-ELSE conditions), one can no longer rely on gradient descent, but instead will need alternative methods such as, e.g., automated program synthesis [7, 14]. Advancements in these fields may be needed before good conversational capabilities can be inferred directly from data.

Furthermore, we also do not exclude the use of black boxes, e.g., in the identification of the user’s intent. However, we are adamant in our belief that black box models should not be used in those parts of the agent that are directly concerned with *decision-making*. If the output of a black box component may have an influence on the decision-making (e.g., when used for identifying user intent), solutions should be studied to identify and minimize the related risks.

5 CONCLUSIONS

To conclude, in only the last few years it has become apparent that conversational AI is reaching a level of maturity such that its use will soon be widespread. Before long, in daily life one might encounter high-quality conversational agents whose output is almost indistinguishable from that of a person. At this crossroads, it is vital to consider the wider implications of conversational AI and to ascertain, to the greatest degree possible, that conversational agents are transparent and safe for use. To this end, we have here presented five key principles that, in our view, should be taken into account in research and applications involving conversational AI. We have also exemplified the application of those principles. In our current work, we are developing an interpretable approach that embodies all five principles, but it is important to note that the principles themselves are more important than any specific implementation.

REFERENCES

- [1] Alberto Bartoli and Eric Medvet. 2020. Exploring the Potential of GPT-2 for Generating Fake Reviews of Research Papers. *Fuzzy Systems and Data Mining VI: Proceedings of FSDM 2020* 331 (2020), 390.
- [2] Rodrigo Bavaresco, Diógenes Silveira, Eduardo Reis, Jorge Barbosa, Rodrigo Righi, Cristiano Costa, Rodolfo Antunes, Marcio Gomes, Clauter Gatti, Mariangela Vanzin, Saint Clair Junior, Elton Silva, and Carlos Moreira. 2020. Conversational agents in business: A systematic literature review and future research directions. *Computer Science Review* 36 (2020), 100239.
- [3] EM Bender, T Gebru, A McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [4] Abeba Birhane and Vinay Uday Prabhu. 2021. Large Image Datasets: A Pyrrhic Win for Computer Vision?. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1536–1546.
- [5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
- [6] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, and Colin Raffel. 2020. Extracting Training Data from Large Language Models. *arXiv preprint arXiv:2012.07805* (2020).
- [7] John Clarke, Jose Javier Dolado, Mark Harman, Rob Hierons, Bryan Jones, Mary Lumkin, Brian Mitchell, Spiros Mancoridis, Kearton Rees, Marc Roper, and Martin J Shepperd. 2003. Reformulating software engineering as a search problem. *IEE Proceedings-software* 150, 3 (2003), 161–175.
- [8] Bob Coecke, Giovanni de Felice, Konstantinos Meichanetzidis, and Alexis Toumi. 2020. Foundations for Near-Term Quantum Natural Language Processing. *arXiv preprint arXiv:2012.03755* (2020).
- [9] Ryan Daws. 2020. Medical chatbot using OpenAI’s GPT-3 told a fake patient to kill themselves. AI News, <https://artificialintelligence-news.com/2020/10/28/medical-chatbot-openai-gpt3-patient-kill-themselves/>. Accessed May 2021.
- [10] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press, Cambridge.
- [11] Sara Hooker. 2020. The hardware lottery. *arXiv preprint arXiv:2009.06489* (2020).
- [12] Anna Jobin, Marcello Lenca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399.
- [13] Daniel Jurafsky and James H. Martin. 2009. *Speech and Language Processing (2nd Edition)*. Prentice-Hall.
- [14] John R Koza. 1992. *Genetic programming: On the programming of computers by means of natural selection*. Vol. 1. MIT press, Cambridge.
- [15] Liliana Laranjo, Adam G Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie Y S Lau, and Enrico Coiera. 2018. Conversational agents in healthcare: A systematic review. *Journal of the American Medical Informatics Association* 25, 9 (2018), 1248–1258.
- [16] David D Luxton. 2020. Ethical implications of conversational agents in global public health. *Bulletin of the World Health Organization* 98, 4 (2020), 285.
- [17] Gary F Marcus and Ernest Davis. 2020. GPT-3, bloviator: OpenAI’s language generator has no idea what it’s talking about. *Technology Review*, <https://www.technologyreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/>. (2020). Accessed May 2021.
- [18] Christoph Molnar. 2020. *Interpretable machine learning*. Lulu.com.
- [19] Daniel W Otter, Julian R Medina, and Jugal K Kalita. 2020. A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems* 32, 2 (2020), 604–624.
- [20] Qing Ping, Feiyang Niu, Govind Thattai, Joel Chengottusseriyil, Qiaozi Gao, Aishwarya Reganti, Prashanth Rajagopal, Gokhan Tur, Dilek Hakkani-Tur, and Prem Nataraja. 2020. Interactive Teaching for Conversational AI. *arXiv preprint arXiv:2012.00958* (2020).
- [21] Hope Reese. 2016. Why Microsoft’s ‘Tay’ AI bot went wrong. *Tech Republic*, <https://www.techrepublic.com/article/why-microsofts-tay-ai-bot-went-wrong/>. Accessed May 2021.
- [22] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (2019), 206–215.
- [23] Walid S Saba. 2007. Language, logic and ontology: Uncovering the structure of commonsense knowledge. *International Journal of Human-Computer Studies* 65, 7 (2007), 610–623.
- [24] Walid S Saba. 2018. Logical semantics and commonsense knowledge: Where did we go wrong, and how to go forward, again. *arXiv preprint arXiv:1808.01741* (2018).
- [25] Mattias Wahde and Marco Virgolin. 2022. Conversational agents: Theory and applications. In *Handbook on Computer Learning and Intelligence (to appear)*. World Scientific.
- [26] Richard S Wallace. 2009. The anatomy of ALICE. In *Parsing the Turing Test*. Springer, 181–210.
- [27] Pengcheng Yin and Graham Neubig. 2017. A Syntactic Neural Model for General-Purpose Code Generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 440–450.
- [28] Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. POMDP-based statistical spoken dialog systems: A review. *Proc. IEEE* 101, 5 (2013), 1160–1179.