Transparency in Language Generation: Levels of Automation

Justin Edwards
University College Dublin
justin.edwards@ucdconnect.ie

Allison Perrone
Bat Camp
allison@batcamp.org

Philip R. Doyle
University College Dublin
philip.doyle1@ucdconnect.ie

ABSTRACT

Language models and conversational systems are growing increasingly advanced, creating outputs that may be mistaken for humans. Consumers may thus be misled by advertising, media reports, or vagueness regarding the role of automation in the production of language. We propose a taxonomy of language automation, based on the SAE levels of driving automation, to establish a shared set of terms for describing automated language. It is our hope that the proposed taxonomy can increase transparency in this rapidly advancing field.

CCS CONCEPTS

• **Human-centered computing** → Natural language interfaces.

KEYWORDS

transparency, natural language generation, language models, taxonomy, levels of automation

ACM Reference Format:

Justin Edwards, Allison Perrone, and Philip R. Doyle. 2020. Transparency in Language Generation: Levels of Automation. In *Proceedings of 2nd Conference on Conversational User Interfaces (CUI '20)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3405755.3406136>

1 INTRODUCTION

Increasingly, transparency is becoming a key practical and ethical concern for developers of natural language technologies [2]. Conversational research suggests knowledge about the type of partner a person is talking to (human or computer) has a significant impact: on language choices in dialogue [5, 6]; on perceptions of a partner's knowledge and capabilities [7]; on perceptions of interpersonal connection [7, 12]; and on perceptions of trustworthiness [16]. Failure to disclose when someone is talking to a computer rather than a person can also lead to heightened expectations about system capability. Subsequent failure to meet these expectations can lead to frustration, limited use and even abandonment of CUIs [8]. Transparency is also a key element of recent policy implementations such as GDPR, which guarantees European citizens a right to transparency in their interactions with technology [11]. Here, we propose a structure for 'levels of automation' that can be used to clearly delineate the roles of humans and machines in generating output. Our hope is that the levels of automation posited here - inspired by the SAE taxonomy of driving automation - will be a

first step toward greater transparency in the field of conversational technology, inspiring iterative refinement of this taxonomy that produces universal descriptions for these technologies.

2 SAE LEVELS

Issued in 2014, the SAE taxonomy of driving automation levels clearly categorizes automated driving systems according to the level of control possessed by the driver and the vehicle, respectively. The clarity of language afforded by the SAE taxonomy allows for marketing of new technologies that avoids misleading consumers [1] as well as allowing researchers, policymakers and journalists to discuss emerging technologies [9]. It also provides consumers with clarity around levels of automation and control. Table 1 presents a summarized version of the SAE levels of driving automation.

Table 1: SAE Levels of driving automation

SAE Level	Description
0	No automation
1	Driver assistance (a single automated system like cruise control)
2	Partial automation (vehicle can operate autonomously, but human monitors the environment and can take control at any time)
3	Conditional automation (vehicle can monitor environment, operate autonomously, but human must be available to takeover in some situations)
4	High automation (under certain circumstances, the vehicle is fully autonomous, human takeover is option in other circumstance)
5	Full automation: (No human interaction required, takeover may be disabled)

SAE levels are defined in terms of the role and responsibility of the driver in relation to the vehicle's automated features. As such, they give insight into varying levels of human/system control across different in-car automation systems and across various driving events a system may or may not be designed to handle. As seen in Table 1, human operators are in control of levels 2 and below, whilst the automated driving system is in control at levels 3 and above. This delineation helps in setting appropriate expectations for drivers, aids policymakers in establishing appropriate legal frameworks, and allows for greater accuracy in reporting when these human-computer interactions fail. The SAE levels were developed through committee discussion among driving engineers and industry stakeholders to develop a common set of terms for this specific issue. It is our hope that by structuring similar levels in the field of language generation, we too can give a common language to our research community, enhancing transparency in the field.

This work is licensed under a Creative Commons Attribution International 4.0 License.

CUI '20, July 22–24, 2020, Bilbao, Spain

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7544-3/20/07.

<https://doi.org/10.1145/3405755.3406136>

3 LEVELS OF LANGUAGE GENERATION AUTOMATION

Like the SAE levels, we propose a structure of 6 levels of automation that define the role of a human author when supported by automated language generation systems, including those using both rule-based and probabilistic approaches. Below, we posit definitions and examples of each of the six levels of language generation.

Level 0: Fully human-written language: indicative of language written and selected exclusively by a human. While editing assistance like spell check may be used at this level, lexical choices are entirely controlled by a human.

Level 1: Language assistance: includes language that is entirely written and selected by a human, but may be scripted to present automatically. These may include highly constrained chatbots designed for a specific role, or phone trees limited by predetermined sequences. Systems like this present users with a limited set of dialogue options per turn, rather than allowing a user to freely enter language. Dialogues may take varied branching paths, but only within the confines of predetermined sets of commands and responses that were generated by a human author. No novel language is generated throughout these interactions. In this way, level 1 automation allows for automation of scripted interactions rather than automation of novel generation.

Level 2: Partial automation: includes language generated through shared effort between a human author and an automated system. This may include language written by a human then selected algorithmically, and/or language generated algorithmically then selected by a human. Many modern Twitterbots take the former approach, including Twitterbots built using Tracery (e.g. Lost Tesla Bot¹). Tracery is a language generation approach that generates text through slot filling using random or conditional selection [4]. An author can create a number of templates, called origins, and define slots and lists of potential entries for each slot. A Twitterbot employing Tracery thus produces novel text by combining several human-written texts.

An example of partial automation can also be seen in output from Voicebox, a predictive text tool by Botnik Studios². Voicebox is a Markov chain-based predictive text keyboard that can be trained on existing text uploaded by the author (e.g., a compilation of an artist's lyrics, a body of text from a novel). The author then generates a new text by selecting one word at a time. Level 2 partial automation still requires a high level of human involvement as language must be initially provided by a human author, but novel generation is possible.

Level 3: Conditional automation: Here language generation is accomplished by automation, with human effort reserved primarily for selecting the generated language for use. At this level, the role of authorship is shared between a human and a machine. As an example, in a recent article in *The Economist*, contributor Tom Standage interviewed GPT-2 by inputting interview questions about technology to watch in 2020 [15]. GPT-2, a large stochastic language model [13], was instantiated in a shared writing environment in which an author could input text cues for GPT-2 to

continue to generate language. Standage authored questions, received several generated answers, and selected which answers to publish. Although a higher level of automation was asserted, this masks the role Standage had in composing the interview, particularly in selecting which answers to use. An interaction like this is more accurately described as a Level 3, conditional automation involving a shared task between a human and a machine.

Level 4: High automation: Level 4 language automation requires no human supervision. Here, language is generated and selected stochastically, though constrained to a specific domain or language task. One currently available implementation of level 4 language automation is AI Dungeon, a GPT-2 based text adventure role-playing game³. Gameplay in AI Dungeon is generated by an instantiation of GPT-2 trained on text from role-playing games, with the game responding to human text input with stochastically generated story development, creating emergent gameplay between the game and player. Similar implementations could be used in a variety of other use-cases, like educational tools to encourage children to write, and marketing tools such as embodied agents promoting brand engagement. While training models and defining domains may require high human effort, language generation is not performed by a human at this level of automation. This fully automatised generation differentiates level 4 from level 3, while domain constraints differentiate it from level 5.

Level 5: Fully automated language: Level 5 represents fully automated language generation. Like level 4, fully automated language requires no human supervision and is both produced and selected stochastically. However, unlike level 4, language is not constrained to a particular domain, task, or topic. This is a long-held goal for general artificial intelligence [10] and has not yet been accomplished. To merit consideration as a level 5 automated language system, a system would need to be capable, without modification, of various language tasks including both task-oriented transactional dialogue and open-ended social dialogue.

Level 5 systems may be most useful as a template that users could then constrain for different specific purposes, thus rendering specific implementations into level 4 systems. While there is some evidence that people enjoy chatting with natural-sounding chatbots [14], other work casts doubt suggesting these preferences may be due to the novelty of the interaction [3], which may mean there are limited use-cases for this level of automation. It should be made clear that the level of automation of a language system does not necessarily correspond to the quality of outputs nor to the utility of the system overall.

4 CONCLUSION

Conversational systems and language generation tools are becoming increasingly advanced, blurring lines between human-generated and computer-generated language. Degrees of automation have been clearly delineated in the field of automated driving though the use of a shared set of definitions that can be understood by a variety of stakeholders. By using a similar taxonomy in the field of conversational technology, we can ensure that this field maintains transparency in discussion of generated language, ensuring consumers are not misled when interacting with these systems.

¹twitter.com/LostTesla

²botnik.org

³aidungeon.io

REFERENCES

- [1] Keith Barry and Jeff Plungis. 2019. Do Levels of Car Automation Matter? *Consumer Reports* (Oct. 2019). <https://www.consumerreports.org/autonomous-driving/levels-of-car-automation/>
- [2] Heloisa Candello, Benjamin Cowan, Cosmin Munteanu, Joel Fischer, Leigh Clark, Stephan Schlogl, Jaisie Sin, Christine Murad, Stuart Reeves, Martin Porcheron, and Chelsea M. Myers. 2020. CUI@CHI: Mapping Grand Challenges for the Conversational User Interface Community. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems - CHI '20*. ACM Press, Honolulu, US. <https://doi.org/10.1145/3334480.3375152>
- [3] Leigh Clark, Cosmin Munteanu, Vincent Wade, Benjamin R. Cowan, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, and Christine Murad. 2019. What Makes a Good Conversation?: Challenges in Designing Truly Conversational Agents. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*. ACM Press, Glasgow, Scotland Uk, 1–12. <https://doi.org/10.1145/3290605.3300705>
- [4] Kate Compton, Ben Kybartas, and Michael Mateas. 2015. Tracery: An Author-Focused Generative Text Tool. In *Interactive Storytelling*, Henrik Schoenau-Fog, Luis Emilio Bruni, Sandy Louchart, and Sarune Baceviciute (Eds.). Springer International Publishing, Cham, 154–161.
- [5] Benjamin R Cowan and Holly P Branigan. 2015. Does Voice Anthropomorphism Affect Lexical Alignment in Speech-Based Human-Computer Dialogue? (2015), 5.
- [6] Benjamin R. Cowan, Philip Doyle, Justin Edwards, Diego Garaialde, Ali Hayes-Brady, Holly P. Branigan, João Cabral, and Leigh Clark. 2019. What's in an accent?: the impact of accented synthetic speech on lexical choice in human-machine dialogue. In *Proceedings of the 1st International Conference on Conversational User Interfaces - CUI '19*. ACM Press, Dublin, Ireland, 1–8. <https://doi.org/10.1145/3342775.3342786>
- [7] Philip R. Doyle, Justin Edwards, Odile Dumbleton, Leigh Clark, and Benjamin R. Cowan. 2019. Mapping Perceptions of Humanness in Intelligent Personal Assistant Interaction. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services - MobileHCI '19*. ACM Press, Taipei, Taiwan, 1–12. <https://doi.org/10.1145/3338286.3340116>
- [8] Ewa Luger and Abigail Sellen. 2016. "Like Having a Really Bad PA": The Gulf between User Expectation and Experience of Conversational Agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. ACM Press, Santa Clara, California, USA, 5286–5297. <https://doi.org/10.1145/2858036.2858288>
- [9] Dimitris Milakis. 2019. Long-term implications of automated vehicles: an introduction. *Transport Reviews* 39, 1 (Jan. 2019), 1–8. <https://doi.org/10.1080/01441647.2019.1545286>
- [10] Nils J. Nilsson. 2010. *The quest for artificial intelligence: a history of ideas and achievements*. Cambridge University Press, Cambridge ; New York. OCLC: ocn396185265.
- [11] European Parliament and Council of the European Union. 2016. *Regulation on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (Data Protection Directive)*. Regulation (EU) 2016/679. European Parliament and Council of the European Union.
- [12] Amanda Purington, Jessie G. Taft, Shruti Sannon, Natalya N. Bazarova, and Samuel Hardman Taylor. 2017. "Alexa is my new BFF": Social Roles, User Satisfaction, and Personification of the Amazon Echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '17*. ACM Press, Denver, Colorado, USA, 2853–2859. <https://doi.org/10.1145/3027063.3053246>
- [13] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019), 9.
- [14] Heung-yeung Shum, Xiao-dong He, and Di Li. 2018. From Eliza to Xiaolce: challenges and opportunities with social chatbots. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (2018), 10–26.
- [15] Tom Standage. 2019. An artificial intelligence predicts the future. *The Economist* (Nov. 2019). <https://worldin.economist.com/edition/2020/article/17521/artificial-intelligence-predicts-future>
- [16] Ilaria Torre, Jeremy Goslin, Laurence White, and Debora Zanatto. 2018. Trust in artificial voices: A congruency effect of first impressions and behavioural experience. In *Proceedings of the Technology, Mind, and Society*. ACM, 40.