



Large language models and the emergence phenomena^{*}

ARTICLE INFO

Keywords

LLMs
ChatGPT
Emergence
Generative models

ABSTRACT

This perspective explores the potential of emergence phenomena in large language models (LLMs) to transform data management and analysis in radiology. We provide a concise explanation of LLMs, define the concept of emergence in machine learning, offer examples of potential applications within the radiology field, and discuss risks and limitations. Our goal is to encourage radiologists to recognize and prepare for the impact this technology may have on radiology and medicine in the near future.

Large language models (LLMs), like ChatGPT, have captured public and scientific attention for their revolutionary impact on natural language processing (NLP) [1,2]. NLP is a technology that converts unstructured human language into structured data, facilitating a wide range of applications. Historically, the field of NLP has seen significant advancements. Early text classification relied on rule-based methods, which, while effective to a degree, were constrained by the need for manual annotation, application limitations, and the inability to comprehend complex linguistic nuances. Over time, these techniques were supplemented and often superseded by machine learning-based methods, and later, deep learning strategies.

In radiology, NLP applications have found utility in voice recognition systems, structuring of radiological reports, information extraction from reports, diagnostic labeling, selecting imaging protocols, and generation of research cohorts [3–5]. In recent years, transformer models have emerged as the state-of-the-art method for many NLP applications due to their superior ability to understand context and manage long texts. LLMs, which are primarily built on the transformer architecture, are currently leading the charge in NLP research and development.

LLMs are trained on large amounts of text data and numerous parameters [6]. According to the scaling laws in machine learning [7], increasing the size of a model generally improves its performance in subsequent NLP tasks. This suggests that these performance gains can be predicted methodologically based on the behaviors observed in smaller, similar models [8,9]. However, in some cases, the performance for some tasks can become unpredictable, and might be described as emergence [8].

Emergence is a term used in different areas, including physics, biology and medicine. It occurs when unpredictable properties arise in large systems, from interactions of simpler components [10,11]. This occurs without explicit programming or design. These new patterns, properties or behaviors are not present in the small components. Thus, they cannot be easily predicted [10].

In medicine, emergence can be seen for example in pandemics. Pandemics arise from interactions between pathogens and humans.

They can be difficult to predict and control, and they can have devastating effects on our society. Another example for emergence phenomenon in medicine relates to cancer. Cancer cells are mutated cells that acquire new properties. The mutations can be unpredictable, and allow the tumor to invade the surrounding tissue, replicate effectively, resist treatments and spread to distant parts of the body.

In the context of artificial intelligence, emergence refers to unpredictable tasks performed by an algorithm, that weren't explicitly programmed into it. One relatively simple example in NLP is word-embedding. Word-embedding is a technique that laid the foundation to deep learning NLP. Vectors are used to represent words numerically in order to analyze relations between them [12]. For example, deducing that the word tarsal relates to the word foot in a similar way that carpal relates to hand [5]. The ability to capture semantic relations between words based on their vectors representation in space can be considered as an emergence phenomenon. That is because such relations were not explicitly pre-programmed and cannot be predicted at all times.

Another more complex example for emergence in image analysis can be generative adversarial networks (GANs). In this case the higher-level behavior arises from the interaction between two neural networks, the generator and the discriminator. The generator creates synthetic images, and the discriminator tries to distinguish between real images, and fake ones that were generated by the generator. Both networks improve their performance through a competitive process. The generator tries to produce more realistic images, and the discriminator tries to better identify the fake images [13]. As a result, they indirectly learn from each other, and their combined performance exhibits emergent behavior. GANs are used in radiology for a variety of tasks. For example, image denoising and enhancement, generating synthetic radiology images for algorithms training, transforming images from one modality to another, and image segmentation [13].

Larger models have more potential for unpredicted behaviors. The ability of LLMs to learn about objects in the world, without being explicitly programmed to recognize them, is one intriguing example. The algorithm can form a concept of a particular object, by identifying

^{*} No funding has been received for this work.

common features in different objects that are labeled similarly [14]. Interactions of simple elements within the algorithm enable it to learn new information, rather than execute pre-programmed instructions. When queried about a chair, ChatGPT can provide extensive information. It was never explicitly taught about this concept, yet it has autonomously learned different features and characteristics of chairs through texts processing. This emergent behavior is particularly fascinating as it depicts the algorithms' ability to autonomously learn about the world around us.

In radiology, we can benefit from emergence in LLMs. Instead of analyzing data based on pre-programmed rules and statistical methods that require large volumes of annotated data, the pre-trained algorithm can improve autonomously with only few domain-specific examples [14]. Furthermore, the algorithms' analysis may be different from the way that humans think. Thus, they may be able to extract information that is obscure to us.

For instance, an LLM can analyze radiology reports and suggest possible differential diagnosis, based on findings and descriptions that it was not explicitly trained on. Furthermore, the algorithm may potentially identify previously unrecognized correlations between patterns descriptions and medical conditions. This emergent capability could enhance diagnoses of complex cases or rare diseases.

Another potential application can be to uncover unknown clinical insights. LLMs can process radiology reports and extract data that may be obscure to humans, such as outcome predictions, or recognition of risk factors for diseases. The algorithm can also process radiology and medical literature, potentially finding meaningful insights that may not be immediately apparent to human readers. This emergent phenomenon could lead to the identification of novel associations in medicine in general, and radiology in particular.

The occurrence of emergent phenomena in AI is expected to increase significantly. As deep learning models, and specifically language models continue to grow, the potential for unpredicted outputs also escalates. Implications of this phenomenon are vast. The current models are likely mere stepping stones to what comes next. In theory, it could pave the way for the creation of more sophisticated systems capable of learning, adapting, and operating autonomously. Some even regard emergence as the cornerstone for the eventual development of artificial general intelligence (AGI) [15].

We are in the early stages of integrating AI models into healthcare practice, in radiology in particular [16]. It's crucial to differentiate between the development of algorithms for research purposes and their practical application in clinical environments. In a research setting, the primary objective is often performance improvement. Yet, when these models are deployed in a clinical setting, additional factors such as limitations, ethical concerns, and inherent risks must be considered [17]. Explainability is one significant limitation. The decision-making processes of deep learning algorithms are often opaque, which can make it challenging to understand how they arrive at a certain output. This lack of transparency may impede error detection, potentially raising reliability concerns. This issue is further amplified in the context of LLMs and emergent phenomena, where unexpected behaviors can complicate the understanding of how inputs are transformed into outputs. This becomes particularly critical in the field of medicine due to its inherent risks.

Emergence is particularly worrisome in cases where the unpredicted output can not be controlled or regulated. For instance, the potential generation of false information, which in the context of clinical application may lead to inappropriate medical decisions and potentially catastrophic consequences. Algorithms can replicate biases and discrimination, based on the data that they were trained on [18]. Another issue is the question of legal responsibility and liability in cases where AI-driven decisions lead to negative outcomes [17]. Other concerns include privacy and security, since they are inevitably vulnerable to cyber-attacks [19]. Regarding the practical limitations of large-scale language models (LLMs) in radiology, substantial computational

resources and energy are necessary for their training. Variables such as the size of the dataset, the model size, and the training volume have all increased dramatically [20]. Additionally, multiple training iterations are often required to minimize random effects and explore the hyper-parameter space, further escalating costs. Finally, the environmental impact of the energy required to train these models should not be overlooked.

To conclude, considering the rapid advancements with LLMs, emergence in AI is likely to expand. There are already numerous examples for emergence in existing algorithms. Radiology, which has consistently been at the forefront of adopting innovative technologies and AI, is expected to maintain this momentum. It is thus imperative for radiologists to be familiar with these concepts. Understanding emergence, radiologists can leverage its potential, while being mindful to challenges and risks associated with this rapidly evolving technology.

CRedit authorship contribution statement

Vera Sorin: Investigation, Writing – original draft. **Eyal Klang:** Conceptualization, Methodology, Resources, Supervision, Writing – review & editing.

Ethical approval details

N/A.

Funding Sources

None.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] X. Yang, A. Chen, N. PourNejatian, et al., A large language model for electronic health records, *npj Digit. Med.* 5 (2022) 1.
- [2] C. Stokel-Walker, ChatGPT listed as author on research papers: many scientists disapprove, *Nature* 613 (7945) (2023) 620–621.
- [3] P. López-Úbeda, T. Martín-Noguerol, K. Juluru, A. Luna, Natural language processing in radiology: update on clinical applications, *J. Am. Coll. Radiol.* 19 (11) (2022) 1271–1285.
- [4] S.C. Fanni, M. Gabelloni, A. Alberich-Bayarri, E. Neri, Structured reporting and artificial intelligence, *Radiologe* (2022) 169–183.
- [5] V. Sorin, Y. Barash, E. Konen, E. Klang, Deep learning for natural language processing in radiology—fundamentals and a systematic review, *J. Am. Coll. Radiol.* 17 (5) (2020) 639–648.
- [6] D. Khurana, A. Koli, K. Khatter, S. Singh, Natural language processing: state of the art, current trends and challenges, *Multimed. Tools Appl.* 82 (3) (2022) 3713–3744.
- [7] J. Kaplan, S. McCandlish, T. Henighan, et al., Scaling laws for neural language models, arXiv preprint arXiv:2001.08361, 2020.
- [8] J. Wei, Y. Tay, R. Bommasani, et al., Emergent abilities of large language models, arXiv preprint arXiv:2206.07682, 2022.
- [9] D. Ganguli, D. Hernandez, L. Lovitt, et al., Predictability and surprise in large generative models, Paper presented at: 2022 ACM Conference on Fairness, Accountability, and Transparency, 2022.
- [10] S. Kivelson, S.A. Kivelson, Defining emergence in physics, *npj Quantum Mater.* 1 (1) (2016).
- [11] T. O'Connor, Emergent properties, 2020.
- [12] V. Sorin, Y. Barash, E. Konen, E. Klang, Deep-learning natural language processing for oncological applications, *Lancet Oncol.* 21 (12) (2020) 1553–1556.
- [13] V. Sorin, Y. Barash, E. Konen, E. Klang, Creating artificial images for radiology applications using generative adversarial networks (GANs) – a systematic review, *Acad. Radiol.* 27 (8) (2020) 1175–1185.
- [14] R. Bommasani, D.A. Hudson, E. Adeli, et al., On the opportunities and risks of foundation models, arXiv preprint arXiv:2108.07258, 2021.
- [15] Se. Bubeck, V. Chandrasekaran, R. Eldan, et al., Sparks of artificial general intelligence: early experiments with GPT-4, *Comput. Lang.* 4 (2023).

- [16] S. Benjamens, P. Dhunoo, B. Meskó, The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database, *npj Digit. Med.* 3 (1) (2020).
- [17] S. Harrer, Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine, *eBioMedicine* 90 (2023), 104512.
- [18] V. Sorin, E. Klang, Artificial intelligence and health care disparities in radiology, *Radiology* 301 (3) (2021) (E443-E443).
- [19] S.G. Finlayson, J.D. Bowers, J. Ito, J.L. Zittrain, A.L. Beam, I.S. Kohane, Adversarial attacks on medical machine learning, *Science* 363 (6433) (2019) 1287–1289.
- [20] O. Sharir, B. Peleg, Y. Shoham, The cost of training nlp models: a concise overview, *arXiv preprint arXiv:2004.08900*, 2020.

Vera Sorin*

Department of Diagnostic Imaging, Chaim Sheba Medical Center, affiliated to the Sackler School of Medicine, Tel-Aviv University, Israel

DeepVision Lab, Sheba Medical Center, Tel Hashomer, Israel

Eyal Klang

*Department of Diagnostic Imaging, Chaim Sheba Medical Center, affiliated to the Sackler School of Medicine, Tel-Aviv University, Israel
DeepVision Lab, Sheba Medical Center, Tel Hashomer, Israel
Sami Sagol AI Hub, ARC, Sheba Medical Center, Israel*

* Correspondence to: Department of Diagnostic Imaging, Chaim Sheba Medical Center, Emek Haela St. 1, Ramat Gan 52621, Israel.
E-mail address: verasrn@gmail.com (V. Sorin).