How can LLMs transform the robotic design process?

Stella, Francesco; Della Santina, Cosimo; Hughes, Josie

DOI
10.1038/s42256-023-00669-7

Publication date
2023

Document Version
Final published version

Published in
Nature Machine Intelligence

Citation (APA)
Stella, F., Della Santina, C., & Hughes, J. (2023). How can LLMs transform the robotic design process? *Nature Machine Intelligence*, *5*(6), 561-564. https://doi.org/10.1038/s42256-023-00669-7

https://doi.org/10.1038/s42256-023-00669-7

How can LLMs transform the robotic design process?

Francesco Stella, Cosimo Della Santina & Josie Hughes

Check for updates

We show that large language models (LLMs), such as ChatGPT, can guide the robot design process, on both the conceptual and technical level, and we propose new human–AI co-design strategies and their societal implications.

Large language models (LLMs)[1] will fundamentally change the robotics landscape by providing robots with the unprecedented capability to understand and analyse natural language. The key advantage of LLMs is their ability to process and internalize large amounts of text data such as instructions, technical manuals and academic articles, and to factually and coherently respond to questions using this implicit knowledge. The potential of utilizing these powerful AI tools within robotics has already been shown through their ability to synthesize code from text prompts[2] and to translate natural-language instructions into actions executable by robots[3,4]. However, recent improvements in the availability and capabilities of LLMs opens new opportunities, and they may now contribute to another bottleneck of robotics, design. Leveraging their emerging capabilities[5,6], LLMs can deliver a dialogue that educates, stimulates and guides humans in building robots. These capabilities could fundamentally change the methodology by which we design robots, changing the role of humans and enriching and simplifying the design process. So, how can LLMs transform the robotic design process, and what are the associated opportunities and risks?

To explore this question, we consider the case study of a human driven by a desire to "help the world with robotics," and we present a robot designed with ChatGPT-3. We approach the task in two steps. In the first, high-level phase, the computer and human collaborate at a conceptual level, discussing ideas and outlining the specifications for the robot design, while in the second phase the physical realization

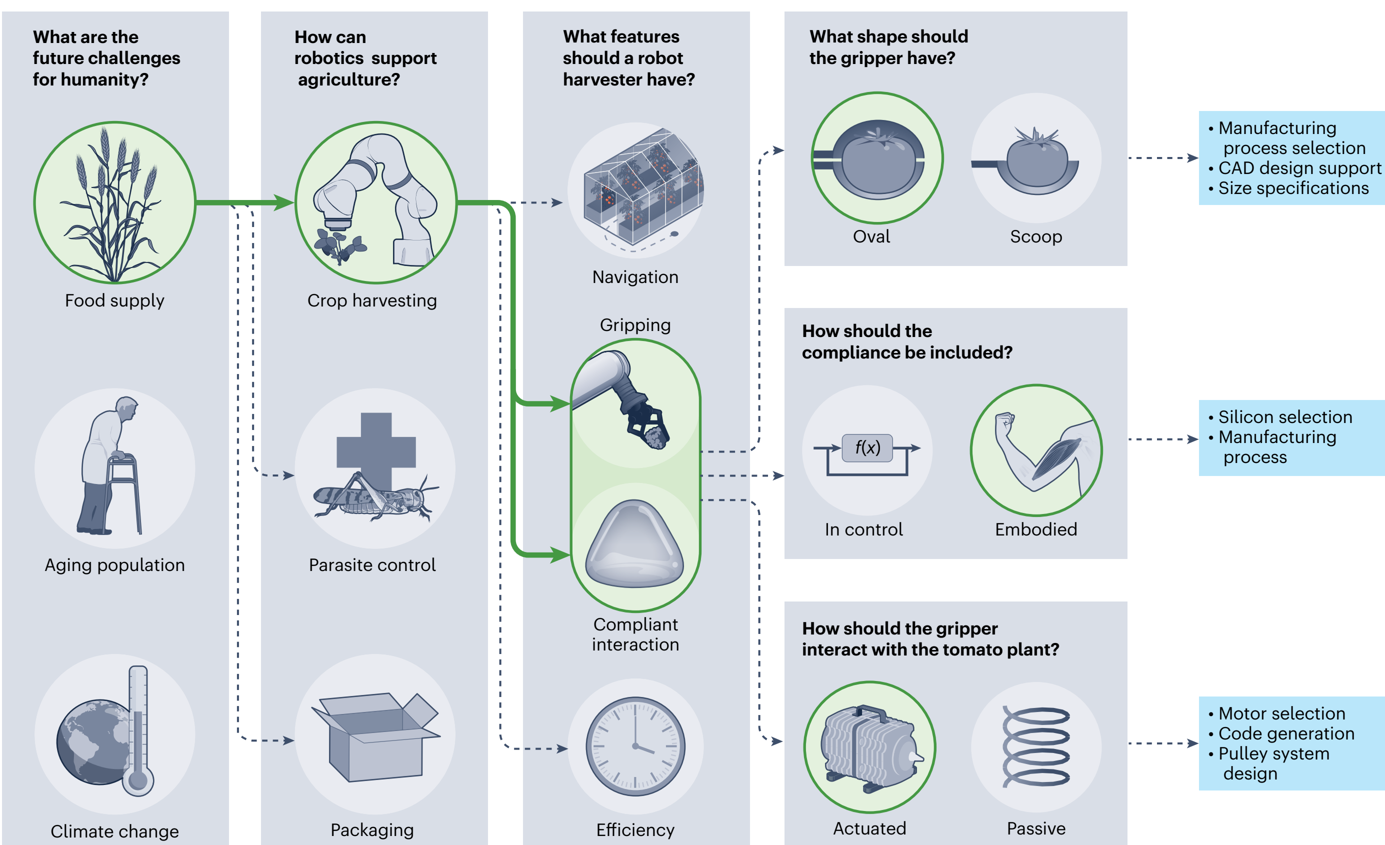

Fig. 1 | Design pipeline. A pictorial overview of the discussion between the human designer and the LLM, with the questions prompted by the human above and the options provided by the LLM below. The green colour highlights the decision tree of the human, who gradually focuses the problem to match their goal.

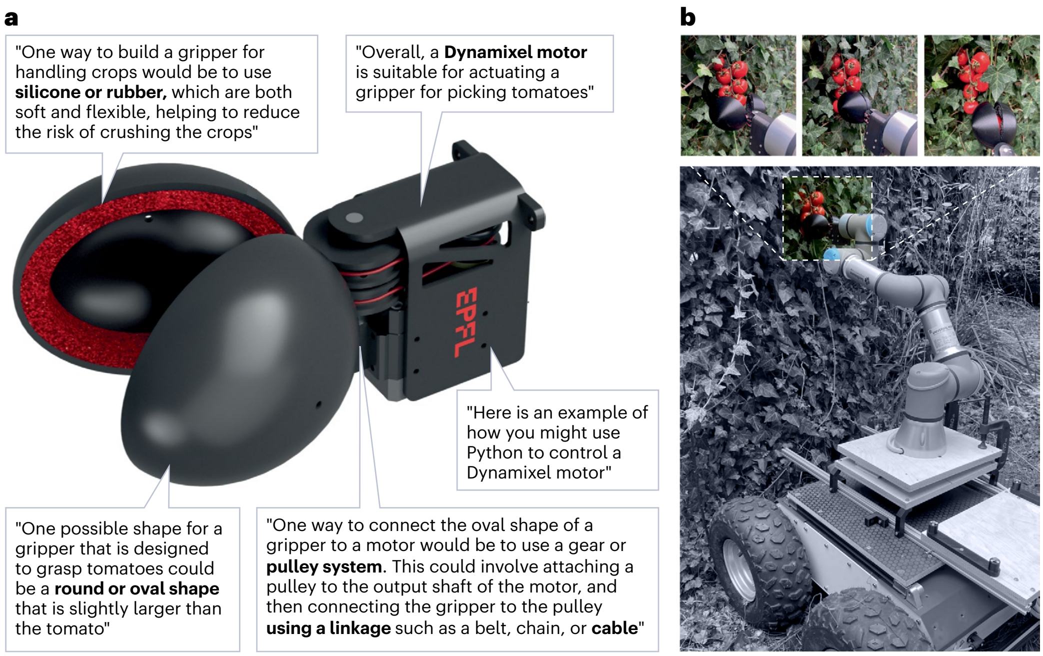

Fig. 2 | An AI model designed this robotic gripper. a, Some of the technical suggestions generated by the LLM, including shape indications, code, component and material selection, and mechanism design. **b**, Guided by these inputs, a gripper was built and tested on real-world tasks, such as tomato picking, as shown at right.

takes place. The full conversation with ChatGPT supporting this case study is provided in the Supplementary Information.

In the ideation phase, the human starts by asking the LLM what the future challenges are for humanity and promptly gets an overview of the main hazards, as shown in Fig. 1. Next, the human selects the most interesting and promising direction, and through further dialogue narrows down the design space. This interaction can span multiple fields of knowledge and levels of abstraction, ranging from concepts to technical implementation. In this process the human is relying on the AI partner to access knowledge outside the human's personal expertise. The AI model helps the human explore intersections between research fields, such as agriculture and robotics, and consider factors that are not part of an engineer's typical training, such as which crop is economically most valuable to automate. Through conversation, the application is selected, and the LLM and the human converge to the technical design specifications, including the software, materials section, mechanism design and manufacturing methods.

In the second, low-level phase of the design process, these directions need to be translated into a physical and functioning robot. Although LLMs cannot currently generate entire CAD models, evaluate code or automatically fabricate a robot, recent advances have shown that AI algorithms can support the technical implementation of software[7], mathematical reasoning[8] or even shape generation[9]. Although we expect that AI approaches will be able to generate these in the future, currently, the technical implementation remains a collaborative effort between AI models and humans. The human takes on a 'technician' role, optimizing the code proposed by the LLM, finalizing the CAD and fabricating the robot. This robot can then be tested in real-world scenarios, and further conversation with the LLM can be used to iterate on the design in the light of experimental evidence. As an example of this second phase, Fig. 2 displays the main outputs generated by the LLM and the real-world deployment of the AI-designed robotic gripper for crop harvesting.

This case study demonstrates the potential for LLMs to transform the design process and how the human–AI relationship may need to vary depending on the expertise of the individual, the stage of the design process and the final goal. By appropriately combining multiple methodologies of human–AI collaboration, the design process can be enhanced and simplified.

At one extreme of human–AI interaction, LLMs could provide all the input required for robot design, which the human follows blindly. The AI is then the inventor, addressing human questions and providing 'creativity', technical knowledge and expertise, whereas the human deals with the technical implementation. This could foster the transfer and democratization of knowledge by enabling nonspecialists to realize robotic systems. For the first time in a computational design framework, the AI agent does not merely solve technical problems specified by the human, but rather proposes conceptual options to the human. In this sense, the LLM acts as the researcher, leveraging knowledge and finding interdisciplinary connections, while the human acts as a manager, providing direction to the design.

A more moderate, yet powerful, approach is a collaborative exploration between the LLM and the human, augmenting the human's expertise by leveraging the ability of the LLM to provide interdisciplinary, wide-ranging knowledge. In this second modality, the role of the LLM is to support the human in efficiently gathering knowledge from fields outside their personal experience, enriching the conception process. For an inherently interdisciplinary research area such as robotics, this has significant potential. By augmenting human knowledge, this methodology removes the limits imposed by a human's education and

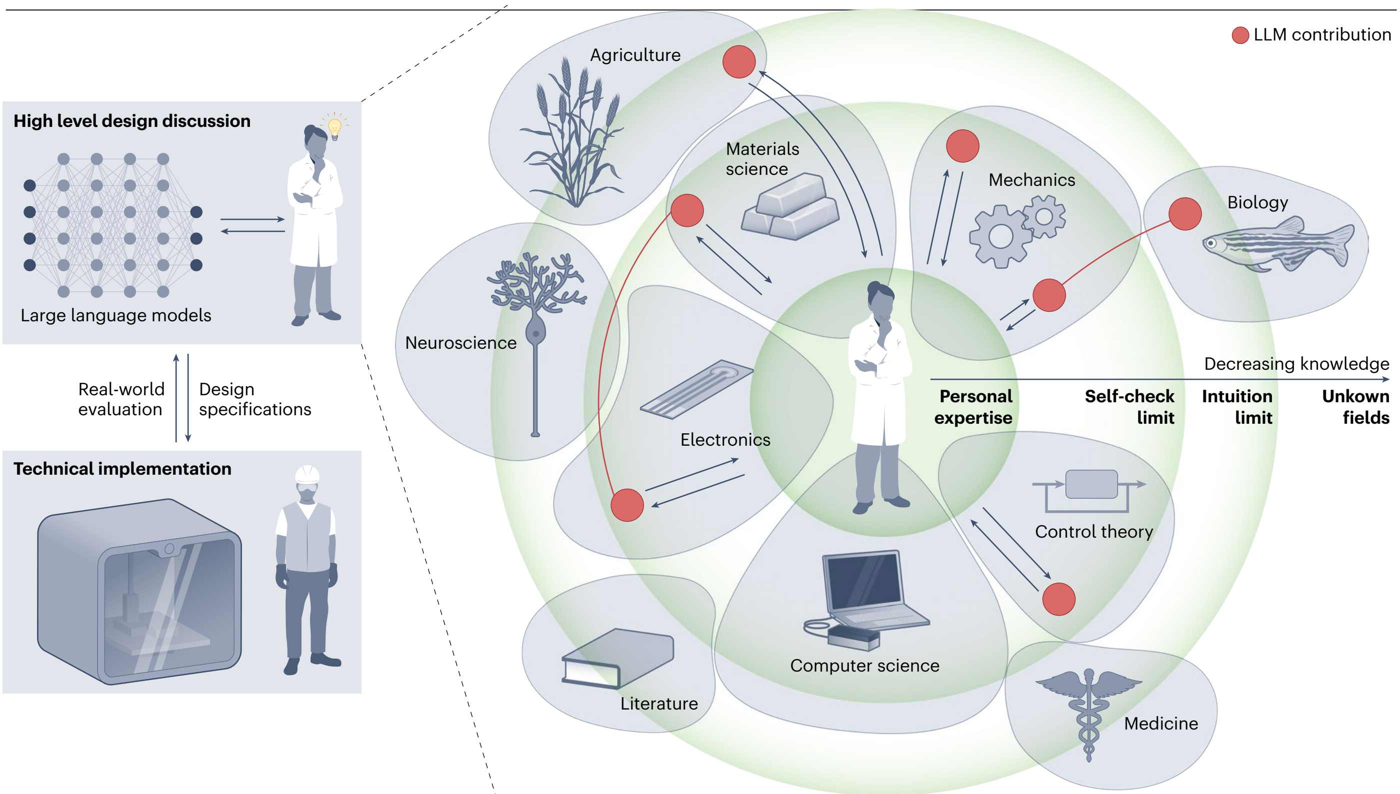

Fig. 3 | Opportunities and risks of human–LLM design collaboration. Left, the two phases of the design process: first the human and LLM discuss the specifics application and of the design, and later the human implements them. Right, a pictorial representation of the knowledge spanned during the high-level discussion. Thanks to the LLM, the human designer can efficiently access zones of knowledge outside their personal expertise and link different of these areas through questions. This, however, comes at the cost of risking accepting wrong inputs for fields too distant from the designer's knowledge. Whereas in a traditional learning process the designer expands their personal experience radially, with LLM-based explorations the designer can access only limited zones of knowledge, thus risking misinterpretation.

supports a human in finding relevant connections between fields, making interdisciplinary research more accessible. However, the knowledge presented by LLMs may be narrow in scope and subject to errors. For fields far from the engineer's expertise, they may not have the ability to fact check the validity of the AI-generated answers. This risk is pictorially represented in Fig. 3. By providing only a small insight or window into vast and complex topics, interactions with LLMs could lead to misinterpretation and oversimplification, ultimately creating errors in the design and biases in the field.

Finally, we can consider a third approach in which the LLM acts as a funnel, helping to refine the design process and providing technical input while the human remains the inventor or scientist involved in the process. The AI can assist with debugging, troubleshooting and the selection of methods, accelerating tedious and time-consuming processes. In this AI–human relationship, the knowledge and intuition of the human moderates the discussion, and the human is working within their scope of expertise so as to be critical of answers and suggestions.

Robotic design is a creative, interdisciplinary and intellectual property (IP)-creating process that currently relies on highly skilled professionals. We believe that a careful combination of these approaches could revolutionize this process. However, introducing LLMs into the design of robots introduces questions regarding potential negative effects. LLMs must be regarded as an evolution of search engines, generating the "most probable" answer to a given prompt[10]. These answers can be incorrect, and when not appropriately fact checked or validated, the LLM output could be misleading or, in the worst case, dangerous. However, unlike search engines, LLMs can propose ways to integrate 'knowledge' and apply it to unseen problems, thus potentially giving a false impression that new knowledge is being generated. This could prevent humans from taking responsibility for the solutions developed[11], which could prohibit and stagnate the advancement of new robotic technologies and designs. Another issue with the widespread use of LLMs in robotic design is the bias toward solutions that the model statistically prefers, which may hinder the exploration of new technological solutions.

Finally, there are key issues related to plagiarism, traceability and IP[12]. Can a design created via LLM be considered to be novel, given that it builds only on prior knowledge[13], and how can this previous knowledge be referenced? As this technology matures, there are also longer-term considerations including data privacy, the frequency of retraining and how new knowledge should be integrated to maintain the usability and relevancy of this tool. There are also significant societal and ethical implications resulting from human–AI interactions for robot design. If LLMs are used to automate high-level cognitive design tasks, humans may instead take on more technical jobs. This could redefine the set of skills that are required by an engineer and their role in the economy and society.

To conclude, the robotics community must identify how to leverage these powerful tools to accelerate the advancement of robots in an ethical, sustainable and socially empowering way. We must develop means of acknowledging the use of LLMs[14] and also of being able to trace the lineage of LLM-aided designs. Looking forward, we strongly believe that LLMs will open up many exciting possibilities and that, if opportunely managed, they will be a force for good. The design process could be fully automated by combining collaborating LLMs to ask and answer questions, with one helping to refine the other. This approach could also be augmented with automated fabrication to allow a fully autonomous pipeline for the creation of bespoke and optimized robotic systems[15]. Ultimately, there is an open question for the future of this field of how these tools can be leveraged to assist robot developers without limiting the creativity, innovation and scientific endeavours required for robotics to rise to the challenges of the twenty-first century.

Francesco Stella [1,2] ✉, **Cosimo Della Santina** [2,3] **& Josie Hughes** [1]
[1]CREATE Lab, EPFL, Lausanne, Switzerland. [2]Department of Cognitive Robotics, Delft University of Technology, Delft, the Netherlands. [3]Institute of Robotics and Mechatronics, German Aerospace Center (DLR), Wessling, Germany.
✉e-mail: francesco.stella@epfl.ch

Published online: 7 June 2023

References

1. Brants, T., Popat, A. C., Xu, P., Och, F. J. & Dean, J. In *Proc. 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)* 858–867 (ACM, 2007).
2. Budzianowski, P. & Vulić, I. Preprint at https://arxiv.org/1907.05774 (2019).
3. Brohan, A. et al. In *Conference on Robot Learning* 287–318 (PMLR, 2023).
4. Huang, W., Abbeel, P., Pathak, D. & Mordatch, I. Preprint at https://arxiv.org/abs/2201.07207 (2022).
5. Wei, J. et al. Preprint at https://doi.org/10.48550/arXiv.2206.07682 (2022).
6. Floridi, L. & Chiriatti, M. *Minds Mach.* **30**(4), 681–694 (2020).
7. Chen, M. et al. Preprint at https://doi.org/10.48550/arXiv.2107.03374 (2021).
8. Wolfram, S. *Stephen Wolfram Writings* https://go.nature.com/3ozHgNf (2023).
9. Ramesh, A. et al. Preprint at https://doi.org/10.48550/arXiv.2204.06125 (2022).
10. Shanahan, M. Preprint at https://doi.org/10.48550/arXiv.2212.03551 (2022).
11. Stokel-Walker, C. *Nature* **613**, 620–621 (2023).
12. Vincent, J. *The Verge* https://go.nature.com/3N2HTIB (2023).
13. George, A. & Walsh, T. *Nat. Mach. Intell.* **4**, 1057–1060 (2022).
14. Anonymous. *Nat. Mach. Intell.* **5**, 1 (2023).
15. Stella, F. & Hughes, J. *Front. Robot. AI* **9**, 1059026 (2023).

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s42256-023-00669-7.

Peer review information *Nature Machine Intelligence* thanks David Howard for their contribution to the peer review of this work.