
How Far Can Advances in NLP Transfer to Biology and Robotics?

Odysseus Kaloidas
UC Berkeley

Abstract

When a method improves natural language processing (NLP), how much of that advance carries over to biology or robotics? Researchers might reuse a model’s design, adapt its training procedure, or draw on knowledge it has already acquired. Each route can help without reproducing the original gain. A robot may understand which object an instruction refers to while still lacking the skill to grasp it. A cell model may recognize a familiar state while missing information needed to predict its response to treatment. These cases make the transfer question concrete: which obstacle did the language advance remove, and does the new task face that same obstacle?

1 An Advance in One Field, a Promise for Others

Suppose a new training method lets a language model learn from half as many examples. Could the same advance halve the data needed to predict a cell’s response to treatment, or the demonstrations needed to teach a robot? This is the promise of transferring advances in natural language processing (NLP), the computational study of human language. But the saving could come from exploiting regularities peculiar to written explanations, or from an improvement in optimization useful almost anywhere. The receiving field has to know which.

There are several ways for progress to travel. Researchers can reuse a model’s design, its training procedure, or knowledge it has already acquired. Multimodal systems add resources learned jointly from text and images. Each can help, but reusing an ingredient does not establish that the original gain will follow it.

2 Three Things That Can Travel

Reusing an architecture changes the machinery available for learning. Vision Transformer adapted the Transformer, a neural-network design that learns relationships among parts of its input, to image recognition [6]. Protein language models learn useful relationships from amino-acid sequences [7]. Neither result requires images or proteins to be statements in a human language.

A training objective can travel separately. Predicting masked elements, deliberately hidden parts of an input, is possible in both text and biological sequences, but the usefulness of that prediction depends on which relationships it encourages the model to learn. A model can become better at reconstructing its input without becoming equally better at every downstream task.

Acquired knowledge travels when a system uses representations or capabilities learned before its current task. This differs from training the same architecture from scratch. The distinction matters when success is attributed to a new algorithm rather than to the resources it inherits.

3 Biology: Recognizing a State and Predicting Its Change

Universal Cell Embedding provides a concrete success. It uses numerical summaries learned from protein sequences within a model of gene expression, the measured activity of genes. It reports useful cell summaries, called embeddings, for datasets and species withheld from cell-model training [8]. Here, knowledge learned from molecular sequences helps organize a different biological measurement. The transfer is substantial without requiring natural-language descriptions of every cell.

Predicting a cell’s response to an intervention raises another problem. Two cells with similar measured expression may respond differently because their histories or environments differ. Observing more examples of the measured state does not necessarily reveal the missing distinction. Dibaieinia and colleagues’ preprint, *Virtual Cells Need Context, Not Just Scale*, argues that contextual coverage governs such transfer [5]. Its point concerns cellular responses, not a universal failure of biological scaling.

Some pretrained representations do improve predictions of responses to genetic changes and chemical treatments over simpler alternatives, as Cole and colleagues report [4]. A separate benchmark finds that rankings depend on the task and evaluation conditions [3]. Both are preprints. The positive result and the variation matter together: prior learning can expose useful biological relationships, but its value depends on which response the model must predict.

This also clarifies why more dimensions are not the whole explanation. A sentence has already selected objects and relations from a richer scene, whereas an assay, a biological measurement procedure, records quantities whose significance remains to be inferred. But compact language can express a difficult problem, and a large measurement can contain exploitable structure. What matters is whether the observations and prior knowledge distinguish the cases that require different answers.

4 Robotics: Understanding More, Moving Better

RT-2 adapts pretrained vision-language models, which process images and text together, to robotic control, transferring web-derived knowledge and improving the use of existing skills in unfamiliar semantic situations [2]. Understanding which object an instruction refers to can make a robot more useful without teaching it a new kind of grasp.

The robot-control model π_0 makes that complementarity explicit. A model trained on images and language is combined with a specialized component that generates continuous sequences of robot movements [1]. Language-related knowledge and motor competence participate in one system, but the latter also depends on robot experience and an appropriate action representation.

Control introduces a further dependency: an action changes the next situation. A poor grasp can move an object beyond reach, so later decisions face conditions absent from the original demonstration. Imitation-learning theory studies this dependence between a policy, the rule for choosing actions, and the states it encounters [9]. Text generation can also drift after its own mistakes. Physical control adds external dynamics, timing constraints, and consequences that cannot always be repaired by generating a better continuation.

An advance in planning may therefore transfer while leaving contact control nearly unchanged. Conversely, an improvement in motor learning can increase task success without improving linguistic reasoning. A receiving domain contains several opportunities for progress; no single language benchmark measures them all.

5 Why Better Search Needs Better Feedback

Spending more computation while answering, called inference or test-time computation, offers a revealing example. Work on increasing this computation in language models shows that allocating computation to search and refinement can improve performance in the studied settings [10]. The transferable idea is to use additional computation to improve a candidate answer. Its benefit depends on how promising candidates are recognized.

Generating more biological hypotheses helps differently when each can be screened cheaply than when distinguishing them requires a lengthy experiment. Planning additional robot trajectories helps

when the simulator captures the relevant contact dynamics; it can reinforce the wrong choice when those dynamics are inaccurate.

Cheaper optimization, more useful prior knowledge, and better search can all produce major gains beyond language. They do not supply missing observations merely by improving computation. Sometimes they make the acquisition of those observations more effective, which is another route by which progress travels.

The robot may understand a new instruction yet still mishandle the object. The cell model may organize familiar states yet fail to predict an unfamiliar intervention. These are places where transfer stops short of the whole task, even when it has already accomplished something useful. Asking which obstacle an NLP advance removes makes its promise for other fields more concrete, with its remaining limits easier to see.

References

- [1] Kevin Black, Noah Brown, Danny Driess, et al. π_0 : A vision-language-action flow model for general robot control. In *Proceedings of Robotics: Science and Systems*, 2025. URL <https://www.roboticsproceedings.org/rss21/p010.html>.
- [2] Anthony Brohan, Noah Brown, Justice Carbajal, et al. RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of the 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 2165–2183, 2023. URL <https://proceedings.mlr.press/v229/zitkovich23a.html>.
- [3] Sally Chen, Roxana Zahedi, Lucy Chhuo, et al. Harmonised benchmarking of foundation models for single-cell and spatial transcriptomics reveals context-dependent generalisation, 2026. URL <https://arxiv.org/abs/2607.17227>. arXiv:2607.17227, preprint.
- [4] Elijah Cole, Geert-Jan Huizing, Sohan Addagudi, et al. Foundation models improve perturbation response prediction, 2026. URL <https://doi.org/10.64898/2026.02.18.706454>. bioRxiv preprint, February 19.
- [5] Payam Dibaeinia, Sudarshan Babu, Mei Knudson, Ali ElSheikh, Yibo Wen, Han Liu, Jason Perera, and Aly A. Khan. Virtual cells need context, not just scale, 2026. URL <https://pmc.ncbi.nlm.nih.gov/articles/PMC12919078/>. bioRxiv preprint, February 9, version 1.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://arxiv.org/abs/2010.11929>.
- [7] Alexander Rives, Joshua Meier, Tom Sercu, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, 2021. doi: 10.1073/pnas.2016239118. URL <https://doi.org/10.1073/pnas.2016239118>.
- [8] Yanay Rosen, Yusuf Roohani, Ayush Agrawal, Leon Samotorčan, Tabula Sapiens Consortium, Stephen R. Quake, and Jure Leskovec. Universal cell embedding provides a foundation model for cell biology. *Nature*, 656:183–191, 2026. doi: 10.1038/s41586-026-10689-z. URL <https://www.nature.com/articles/s41586-026-10689-z>.
- [9] Stéphane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 627–635, 2011. URL <https://proceedings.mlr.press/v15/ross11a.html>.
- [10] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024. URL <https://arxiv.org/abs/2408.03314>. arXiv:2408.03314. Preprint version consulted.