

1 Foundational AI could usher in a new era for models of all life on Earth

2
3 Joseph Millard^{1*}, Arman Pili¹, Katherine Berthon¹, David A. Coomes², Srinivasan Keshav³, Anil
4 Madhavapeddy³, Robert J. Fletcher¹, Lynn V. Dicks¹

5
6 ¹Department of Zoology, University of Cambridge, Cambridge, UK

7 ²Department of Plant Sciences, Conservation Research Institute, University of Cambridge,
8 Cambridge, UK

9 ³Department of Computer Science and Technology, University of Cambridge, William Gates
10 Building, 15 JJ Thomson Avenue, Cambridge, United Kingdom

11
12 *To whom correspondence should be addressed: jwm57@cam.ac.uk

13
14 *Deep-learned foundational AI could usher in a new era for the simulation of whole ecosystems,*
15 *argue Joseph Millard and colleagues.*

16
17 “Time to model all life on Earth”¹. It’s been just over a decade now since Drew Purves and
18 colleagues wrote this sentence, challenging ecologists to develop models for simulating whole
19 ecosystems. The authors then introduced the Madingley model¹, a new form of rule-based general
20 ecosystem model built for the simulation of both terrestrial and marine systems. Taking inspiration
21 from general circulation models and the climate, Purves and colleagues argued that to fully
22 understand how ecosystems will respond to human activities and environmental change,
23 simulation of these ecosystems would be key. This is likely still true. Process-based models such
24 as Madingley can theoretically predict events that haven’t yet occurred, but might do under future
25 conditions.

26
27 Developed bottom-up from the flow of carbon, plants in the Madingley model were represented
28 via spatially explicit rules for primary production and biomass allocation, driven by environmental
29 inputs such as temperature, precipitation, and remotely sensed net primary productivity. Higher
30 trophic levels were grouped on the basis of traits such as diet, metabolism, and body size, with
31 interactions parameterised on the basis of allometric and metabolic scaling laws². Grouping
32 organisms in this way allowed simulation all the way from plankton and aphids up to whales and
33 elephants.

34
35 The Madingley model represented a significant step-forward for ecosystem modelling. But
36 although highly innovative, the grand challenge of reasonably simulating whole ecosystems has
37 never been fully resolved. Whilst general ecosystem models have advanced since the Madingley
38 model³, they still face fundamental limitations. Practically they are difficult to parameterise, apply,
39 and validate at scales such that they are useful for ecology and conservation. Conceptually, it
40 remains difficult to represent important ecological processes such as organismal intelligence, and
41 to simulate across orders of magnitude in scales of time, space or organismal size⁴. We think that
42 the simulation of ecosystems may be about to take another step forwards, but using a different
43 approach. It’s possible that foundational AI elevates our ability to simulate ecosystems.

44 *A new philosophy on simulation*

45
46 Since the development of Madingley in 2013, neural networks have emerged as a new paradigm
47 in the prediction of natural phenomena. Deep-learning (i.e., a multilayered neural network) has
48 shown that where clean training data and compute are easily available, these models can
49 significantly outperform equivalent physical rule-based algorithms. This has been true for protein
50

51 folding⁵. It has been true for complex video games⁶. It has been true for weather prediction⁷. In
52 each of these domains, deep-learning has raised the level of what's tractable for simulation.

53
54 For ecology and conservation, deep-learning has revolutionised prediction in individual domains
55 such as visual and audio recognition⁸, with tagged datasets facilitating supervised learning. But
56 for the simulation of ecosystems, it's currently not clear how a deep-learning model might be
57 trained, or even what it might look like if we did. We think there might be a way. How? A solution
58 might come from repurposing and combining discrete forms of foundational AI.

59 *On the shoulders of scale*

60
61
62 In 2020, neural network scaling laws arrived⁹. A team of researchers at OpenAI demonstrated
63 that through scaling-up 3 parameters—neural network model size, dataset size, and the quantity
64 of compute—large language models (LLM), a class of self-supervised learning models, get
65 predictably better at predicting the next word⁹.

66
67 ChatGPT was a bet on the philosophy of scale. Build a very large neural network for predicting
68 the next word, train it on a large amount of text and compute, and hope that LLM scaling continues
69 to hold.

70
71 The bet paid off. The eventual result was an LLM that could not only use language in a human-
72 like manner, but also appeared to gain capabilities across many novel tasks⁸. We now call this
73 form of model “foundational AI”—a pretrained, self-supervised general-purpose neural network,
74 on which further fine-tune improvements can be made—and it is rapidly becoming ubiquitous in
75 many parts of society, lifting the level on what an individual can do with a computer. Scaling has
76 since been applied across multiple other domains, with new forms of foundational AI appearing
77 across multiple modalities (e.g., speech, text, images). Foundational AI models require no expert
78 labelling, and thus can exploit large, unlabeled, noisy observational data. Ecologists now use
79 foundational AI models to accelerate their analytical workflows¹¹. Applications for the development
80 of ecological models themselves are growing¹².

81 *Workshop a way forward*

82
83
84 As part of an ARIA (Advanced Research + Invention Agency) funded pre-programme discovery
85 project, over four days we invited world-leading experts in computer science and ecology to two
86 workshops in Cambridge, UK. Our aim was to explore the potential of deep-learned foundational
87 AI in ecology, particularly with respect to the forecasting of ecosystem resilience and the
88 development of a new type of general ecosystem model. We focussed on two types of
89 foundational AI: multi-modal Earth-observation models and generative agent-based simulation.
90 In the context of ecosystem modelling, our reasoning was that Earth-observation models might
91 help parameterize the dynamics of vegetation growth, whereas generative agent-based models
92 might better enable the simulation of animal populations.

93
94 A multi-modal Earth-observation model uses self-supervised learning to compress large
95 quantities of Earth observation data into vector embeddings. These embeddings embody
96 significant information about the structure of the Earth's surface, often including multiple data
97 modalities, in a form that's both repurposable and easy to handle. Two examples are TESSERA¹³
98 and AlphaEarth¹⁴, respectively developed by the University of Cambridge and Google Deepmind.
99 Where some form of large-scale spatially explicit task is concerned, multi-modal Earth-
100 observation models will likely come to form an initial building block across many analytical
101 pipelines.

102
103 Generative agent-based simulation is a new form of agent-based model, in which behaviour of
104 individual agents emerges from multiple generative AI models running simultaneously. Concordia
105 is one example platform¹⁵, inspired by the game Dungeons and Dragons. Developed by Google
106 Deepmind, in Concordia each agent is an instantiation of an LLM, hosted such that it can interact
107 with multiple other LLMs. One LLM agent is the 'Game Master', responsible for tracking the
108 environment at any given time and retaining a history of all prior interactions. All other LLM agents
109 are characters in the simulation, playing out their roles in real-time. Crucially, Concordia assumes
110 that LLM agents, given their training, can act reasonably and apply common sense in any given
111 circumstance.

112
113 Since LLMs are trained on human data, most work using generative agent-based models has
114 concerned the simulation of human social behaviours. For example, generative agent-based
115 simulation has been used to show that strongly-held views on COVID-19 were underpinned by
116 an ideological landscape in existence before the pandemic¹⁶. But if LLMs contain significant
117 ecological knowledge, and they can reason about how this knowledge relates to behaviour, then
118 it's possible generative agent-based simulation will also be applicable in ecology. Promisingly,
119 recent research has shown that the model weights alone (i.e., without consulting the web) of an
120 LLM do indeed possess a significant degree of ecological knowledge¹⁷.

121
122 Over the course of our programme discovery project and workshop, a broad idea emerged: If the
123 embedding of a multi-modal Earth-observation model can be fed to an environment tracking LLM,
124 and other LLMs can interpret that embedding reasonably, then perhaps there could be a way to
125 simulate specific ecological interactions, for any specific location on Earth's surface. The benefit
126 of a modelling platform such as Concordia is that the environment in which agents interact can
127 be parameterised, held constant, and tracked by the Game Master. What this means in the
128 context of an ecosystem model, is that the ecological phenomena of interest can be teased apart
129 from the environment in which it exists. That might mean we could simulate the spread of the
130 reintroduction of beavers, for example, without needing to explicitly parameterise their
131 environment. In other words, the fabric and nuance of ecological and environmental reality would
132 be learned into a foundational neural network, and then generated on-the-fly by the Game Master
133 only as it's encountered in simulation.

134
135 *The momentum is there*

136
137 Generative agent-based general ecosystem simulation could be fruitful, but there is a problem
138 with this type of approach. For any simulation, quantifying its validity is challenging. Without a
139 held-out empirical truth with which to validate, it's not clear any prediction of a generative agent
140 is reasonable. This will be true even where a model is parameterised via a multi-modal Earth-
141 observation embedding.

142
143 Growing data sources might help solve this problem. Since the release of Madingley, total
144 observations on iNaturalist have increased exponentially, now numbering 200 million unique
145 observations¹⁸. Standardised and large databases of biodiversity data accumulated from the
146 literature abound^{19,20}, further accelerated by the accumulation of evidence via automated
147 approaches^{22,23}. And momentum has grown for multi-scale standardised biodiversity monitoring²².
148 Moving forward we must anticipate the ground-truth data we need, and then collect it deliberately.
149 Growing biodiversity data should be held out from the training data underpinning foundational AI
150 simulations, at least in part. Held-out data can then be used to post-train or validate the predictions
151 of simulation, helping to ensure models are not just memorising, allowing a better test of

152 generalisability. It's also possible that as data scales, we find new ways to directly train
153 foundational AI ecological models, perhaps without needing to simulate generative agents at all²⁴.

154
155 The Madingley model represented an ambitious step forward, but ecosystem models have not
156 yet gained traction with conservation decision makers. Conventional general ecosystem models
157 are based on a set of bottom-up governing rules that are challenging to parameterize,
158 contextualise precisely, and validate. We think this may now change. We think it's possible
159 ecological simulation in the future might be learned rather than governed. Generative agent-based
160 models and multi-modal Earth-observation models might be the way. Foundational AI could usher
161 in a new era for the simulation of whole ecosystems.

162

163

164 **Acknowledgements**

165

166 Thanks to all ARIA workshop attendees across our two workshops held at Pembroke College,
167 entitled "Foundational AI to forecast ecosystem resilience". Thanks especially to Mike Harfoot for
168 providing comments before submission of this manuscript. This work was funded by ARIA under
169 one of their Engineering Ecosystem Resilience pre-programme discovery projects. JM is funded
170 by the Leverhulme Trust and the Isaac Newton Trust on an Early Career Fellowship.

171

172 **References**

173

174 1. Purves, D., Scharlemann, J. P. W., Harfoot, M., Newbold, T., Tittensor, D. P., Hutton, J.,
175 & Emmott, S. (2013). Time to model all life on Earth. *Nature*, 493(7432), 295–297.
176 <https://doi.org/10.1038/493295a>

177 2. Brown, J. H., Gillooly, J. F., Allen, A. P., Savage, V. M., & West, G. B. (2004). Toward a
178 metabolic theory of ecology. *Ecology*, 85(7), 1771–1789. <https://doi.org/10.1890/03-9000>

179 3. Töpper, J. P., Hoeks, S., Hartig, F., Kolstad, A., Tucker, M. A., Bartlett, J., Bærum, K. M.,
180 Chipperfield, J., Harfoot, M., Huijbregts, M. a. J., & Svenning, J. (2025). Advancing
181 general ecosystem models (GEMs): Towards a mechanistic understanding of the
182 biosphere in the light of the Anthropocene. *Ecological Solutions and Evidence*, 6(2).
183 <https://doi.org/10.1002/2688-8319.70053>

184 4. Groner, V. P., Cook, J., Orme, C. D. L., Amarasekare, P., Comyn-Platt, E., Rallings, T.,
185 Joshi, J., & Ewers, R. M. (2025). Harmonizing nature's timescales in ecosystem models.
186 *Trends in Ecology & Evolution*, 40(6), 575–585.
187 <https://doi.org/10.1016/j.tree.2025.03.011>

188 5. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O.,
189 Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C.,
190 Kohl, S. a. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler,
191 J., ... Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold.
192 *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>

193 6. Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi,
194 D. H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I.,
195 Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., ... Silver, D. (2019).
196 Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*,
197 575(7782), 350–354. <https://doi.org/10.1038/s41586-019-1724-z>

198 7. Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D.,
199 Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., Lam, R., & Willson, M. (2024).
200 Probabilistic weather forecasting with machine learning. *Nature*, 637(8044), 84–90.
201 <https://doi.org/10.1038/s41586-024-08252-9>

- 202 8. Borowiec, M. L., Dikow, R. B., Frandsen, P. B., McKeeken, A., Valentini, G., & White, A.
203 E. (2022). Deep learning as a tool for ecology and evolution. *Methods in Ecology and*
204 *Evolution*, 13(8), 1640–1660. <https://doi.org/10.1111/2041-210x.13901>
- 205 9. Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S.,
206 Radford, A., Wu, J., & Amodei, D. (2020). Scaling laws for neural language models.
207 arXiv. <https://doi.org/10.48550/arxiv.2001.08361>
- 208 10. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P.,
209 Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023).
210 Sparks of Artificial General Intelligence: Early experiments with GPT-4. arXiv.
211 <https://doi.org/10.48550/arxiv.2303.12712>
- 212 11. Cooper, N., Clark, A. T., Lecomte, N., Qiao, H., & Ellison, A. M. (2024). Harnessing large
213 language models for coding, teaching and inclusion to empower research in ecology and
214 evolution. *Methods in Ecology and Evolution*, 15(10), 1757–1763.
215 <https://doi.org/10.1111/2041-210x.14325>
- 216 12. Rafiq, K., Beery, S., Palmer, M. S., Harchaoui, Z., & Abrahms, B. (2025). Generative AI
217 as a tool to accelerate the field of ecology. *Nature Ecology & Evolution*, 9(3), 378–385.
218 <https://doi.org/10.1038/s41559-024-02623-1>
- 219 13. Feng, Z., Atzberger, C., Jaffer, S., Knezevic, J., Sormunen, S., Young, R., Lisaius, M.C.,
220 Immitzer, M., Jackson, T., Ball, J., Coomes, D.A., Madhavapeddy, A., Blake, A., &
221 Keshav, S. (2025). TESSERA: Precomputed FAIR Global Pixel Embeddings for Earth
222 Representation and Analysis. arXiv. <https://arxiv.org/abs/2506.20380>
- 223 14. Brown, C.F., Kazmierski, M.R., Pasquarella, V.J., Rucklidge, W.J., Samsikova, M.,
224 Zhang, C., Shelhamer, E., Lahera, E., Wiles, O., Ilyushchenko, S. and Gorelick, N., ...
225 Kohli, P. (2025). Alphaearth foundations: An embedding field model for accurate and
226 efficient global mapping from sparse label data. arXiv.
227 <https://doi.org/10.48550/arXiv.2507.22291>
- 228 15. Vezhnevets, A.S., Agapiou, J.P., Aharon, A., Ziv, R., Matyas, J., Duéñez-Guzmán, E.A.,
229 Cunningham, W.A., Osindero, S., Karmon, D. and Leibo, J.Z., (2023). Generative agent-
230 based modeling with actions grounded in physical, social, or digital space using
231 Concordia. arXiv. <https://doi.org/10.48550/arXiv.2312.03664>
- 232 16. Kozłowski, A. C., Kwon, H., & Evans, J. A. (2024). In Silico Sociology: Forecasting
233 COVID-19 Polarization with Large Language Models. arXiv.
234 <https://arxiv.org/abs/2407.11190>
- 235 17. Dorm, F., Millard, J., Purves, D., Harfoot, M. and Mac Aodha, O., (2026). Large language
236 models possess some ecological knowledge, but how much? *Ecological Informatics*.
237 95(103699). <https://doi.org/10.1016/j.ecoinf.2026.103699>
- 238 18. Mason, B. M., Mesaglio, T., Heitmann, J. B., Chandler, M., Chowdhury, S., Gorta, S. B.
239 Z., Grattarola, F., Groom, Q., Hitchcock, C., Hoskins, L., Lowe, S. K., Marquis, M.,
240 Pernat, N., Shirey, V., Baasanmunkh, S., & Callaghan, C. T. (2025). iNaturalist
241 accelerates biodiversity research. *BioScience*, 75(11), 953–965.
242 <https://doi.org/10.1093/biosci/biaf104>
- 243 19. Hudson, L. N., Newbold, T., Contu, S., Hill, S. L. L., Lysenko, I., De Palma, A., Phillips,
244 H. R. P., Alhousseini, T. I., Bedford, F. E., Bennett, D. J., Booth, H., Burton, V. J., Chng,
245 C. W. T., Choimes, A., Correia, D. L. P., Day, J., Echeverría-Londoño, S., Emerson, S.
246 R., Gao, D., . . . Purvis, A. (2016). The database of the PREDICTS (Projecting
247 Responses of Ecological Diversity In Changing Terrestrial Systems) project. *Ecology and*
248 *Evolution*, 7(1), 145–188. <https://doi.org/10.1002/ece3.2579>
- 249 20. Dornelas, M., Antão, L. H., Bates, A. E., Brambilla, V., Chase, J. M., Chow, C. F. Y.,
250 Fontrodona-Eslava, A., Magurran, A. E., Martins, I. S., Moyes, F., Sagouis, A., Adu-
251 Acheampong, S., Acquah-Lampsey, D., Adam, D., Ajani, P. A., Albaina, A., Almaraz, P.,
252 An, J., Anderson, R. S., . . . Zvuloni, A. (2025). BioTIME 2.0: Expanding and improving a

- 253 database of biodiversity time series. *Global Ecology and Biogeography*, 34(5).
254 <https://doi.org/10.1111/geb.70003>
- 255 21. Scheepens, D., Millard, J., Farrell, M., & Newbold, T. (2024). Large language models
256 help facilitate the automated synthesis of information on potential pest controllers.
257 *Methods in Ecology and Evolution*, 15(7), 1261–1273. [https://doi.org/10.1111/2041-](https://doi.org/10.1111/2041-210x.14341)
258 [210x.14341](https://doi.org/10.1111/2041-210x.14341)
- 259 22. Gonzalez, A., Vihervaara, P., Balvanera, P., Bates, A. E., Bayraktarov, E., Bellingham,
260 P. J., Bruder, A., Campbell, J., Catchen, M. D., Cavender-Bares, J., Chase, J., Coops,
261 N., Costello, M. J., Czúcz, B., Delavaud, A., Dornelas, M., Dubois, G., Duffy, E. J.,
262 Eggermont, H., . . . Torrelio, C. Z. (2023). A global biodiversity observing system to unite
263 monitoring and guide action. *Nature Ecology & Evolution*, 7(12), 1947–1952.
264 <https://doi.org/10.1038/s41559-023-02171-0>
- 265 23. Besson, M., Alison, J., Bjerge, K., Goroehowski, T.E., Høye, T.T., Jucker, T., Mann, H.M.
266 and Clements, C.F., 2022. Towards the fully automated monitoring of ecological
267 communities. *Ecology letters*, 25(12), pp.2753-2775. <https://doi.org/10.1111/ele.14123>
- 268 24. Gardiner, R. et al. (2026). Towards ecologically meaningful foundation models.
269 *EcoEvoRxiv*. <https://doi.org/10.32942/X2VM1Z>
270