nature ecology & evolution

Progress

https://doi.org/10.1038/s41559-024-02623-1

Generative AI as a tool to accelerate the field of ecology

Received: 21 July 2023

Accepted: 6 December 2024

Published online: 29 January 2025

Check for updates

Kasim Rafiq[1], Sara Beery[2], Meredith S. Palmer[3], Zaid Harchaoui[4] & Briana Abrahms[1]

The emergence of generative artificial intelligence (AI) models specializing in the generation of new data with the statistical patterns and properties of the data upon which the models were trained has profoundly influenced a range of academic disciplines, industry and public discourse. Combined with the vast amounts of diverse data now available to ecologists, from genetic sequences to remotely sensed animal tracks, generative AI presents enormous potential applications within ecology. Here we draw upon a range of fields to discuss unique potential applications in which generative AI could accelerate the field of ecology, including augmenting data-scarce datasets, extending observations of ecological patterns and increasing the accessibility of ecological data. We also highlight key challenges, risks and considerations when using generative AI within ecology, such as privacy risks, model biases and environmental effects. Ultimately, the future of generative AI in ecology lies in the development of robust interdisciplinary collaborations between ecologists and computer scientists. Such partnerships will be important for embedding ecological knowledge within AI, leading to more ecologically meaningful and relevant models. This will be critical for leveraging the power of generative AI to drive ecological insights into species across the globe.

Over the past decade, the application of AI to ecology has revolutionized our ability to understand the mechanisms governing the structure of ecological systems and the drivers of evolutionary change[1,2]. From enabling the rapid and cost-effective processing of environmental big data for ecosystem monitoring[3] to quantifying species interactions in ecological networks[4], AI has swiftly gained traction as a tool to provide deeper insights into ecology and evolution across diverse taxa and systems.

AI's initial impact within science, including ecology, has been predominantly through AI methods wherein models learn to make predictions or classifications from training data that generalize well to previously unseen data from the same distributions (terms are further defined in Table 1). It has enabled the prediction of protein structures[5], the classification of wildlife and plant species within images[3] and the optimization of traffic flow in urban areas[6]. As the AI discipline has evolved, it has progressed to include a class of AI algorithms capable of generating new data with the statistical patterns and properties of the data upon which the models were trained, herein referred to as generative AI. Though some downstream applications of non-generative and generative AI can overlap, such as in classification and prediction tasks, the algorithms used to learn from the training data and their subsequent capacity to generate realistic data are what define generative AI models (Box 1). Prominent examples of generative AI models include GPT-4o[7], which can generate text, Stable Diffusion[8], which can generate images, and Runway's Gen models[9], which can generate video. Moreover, in addition to these large, computationally

[1]Center for Ecosystem Sentinels, Department of Biology, University of Washington, Seattle, WA, USA. [2]AI and Decision Making, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. [3]Center for Biodiversity and Global Change, Yale University, New Haven, CT, USA. [4]Allen School in Computer Science and Engineering, University of Washington, Seattle, WA, USA.
e-mail: kasim.rafiq@hotmail.co.uk

 https://doi.org/10.1038/s41559-024-02623-1

Table 1 | Common generative AI terminology and definitions for ecologists

Term	Definition
Bias	In the context of generative AI, these are systematic errors or misrepresentations that arise from patterns in the training data that favour certain groups or ideas, perpetuate stereotypes or lead to incorrect assumptions.
Data augmentation	The process of generating new synthetic data to improve an existing dataset of real data.
Diffusion models	A type of generative AI model that uses a series of forward (noise-adding) and reverse (noise-removing) steps to gradually learn the underlying structures of the data for generation.
Fine tuning	The process of further training a pre-trained model to improve its performance in a specific task or domain. For example, foundation large language models can be fine tuned for academic writing applications by providing additional academic texts for training.
Foundation models	AI models trained on large quantities of training data that enable them to carry out a broad range of tasks. For example, foundation large language models can be used to summarize, label and translate texts. These models can be fine tuned for specific tasks or domains.
Generative adversarial networks	A type of generative AI model that learns to generate outputs by generating new samples, evaluating them against real data and penalizing inaccuracies to improve subsequent generations.
Generative AI	A class of AI models that generate new data with the statistical patterns and properties of the data upon which the models were trained.
Hallucinate/hallucination	The generation of false or misleading information not supported by the input data or training dataset. It often occurs when the model generates plausible but factually incorrect or nonsensical outputs.
Large language models	A type of generative AI model trained to understand, interpret and generate language. They are trained on vast text datasets and can perform various language-related tasks.
Multimodal AI	AI (generative or non-generative) that integrates multiple data types, such as text, audio and image data, to perform tasks, including data generation, prediction and classification.
Retrieval-augmented generation	A technique for improving the ability of large language models to retrieve information from external sources that the models were not necessarily trained on.
Training data	In the context of AI, a dataset used to train AI models.
Unsupervised learning	A method for training AI models where the model learns from training data that have not been pre-labelled or annotated.

resource-intensive models, generative AI models are increasingly being developed that are capable of running on consumer hardware and by a greater diversity of independent research teams[10]. To date, generative AI models have been used to generate a variety of data types, from bio-climatic variables for improved weather forecasting[11] to genetic sequences for creating new protein structures[12].

Applications of generative AI in ecology

In this Progress, we aim to highlight the wide range of potential opportunities and challenges of applying generative AI as a tool in ecology. Our goal is not to rank or endorse generative AI methods against non-generative AI, traditional statistical inference or other established methods within the field, as the advantages of different approaches will often be context-dependent, and in many cases there remains uncertainty on whether and how generative AI will affect ecology. Instead, we reference these other approaches to contextualize the potential of this rapidly evolving technology alongside familiar methods. By doing so, we aim to encourage discussion on the broad diversity and scope of generative AI's current and future role within the discipline rather than focussing only on established use cases. We focus primarily on applications unique to ecology and discuss a selection of (non-exhaustive) potential use cases, such as using generative AI models to enrich data-scarce ecological datasets, extend observations of ecological patterns, build upon existing mechanistic modelling frameworks and increase the accessibility of ecological data. Additional potential applications of generative AI within ecology and suggested references for further reading can be found in Table 2. We close by summarizing key considerations and challenges for using these models within ecology.

Augmenting data-scarce datasets

In recent years, technological advances have greatly enhanced data collection in ecology, yielding vast volumes of diverse data, from remotely sensed satellite images to animal tracking data to species detections from autonomous camera trap arrays[13,14]. However, the increased scale of incoming data has brought associated challenges in data processing, resulting in increased times for deriving scientific insights. These challenges, alongside the increasing accessibility of expert-annotated datasets, have catalysed interdisciplinary collaborations between ecologists and computer scientists to develop AI models, typically non-generative, capable of rapidly and accurately classifying and processing data, such as identifying species in camera trap images and plants in citizen-science sightings[3,15,16]. This has saved enormous amounts of time and money for ecology and conservation[15]. For AI models to perform accurately across different conditions, such as different ecosystems, they need to be trained on a large number of representative samples of every classification category across all contexts where the model will be used. However, expert-annotated datasets rarely cover all classification categories and environments uniformly. For example, in the context of camera trap images, data for rare species are less abundant than for common ones, and models are often trained in specific ecosystems that are not generalizable across the entire focal species' range[17,18]. Particularly for wildlife species known to be elusive with large and diverse ranges, which encompasses many threatened taxa, or rare endemic plant species, these data may be impossible to collect in the wild.

Data augmentation is a tactic used to supplement real data for training AI models. Datasets can be augmented to increase the amount and diversity of training data through methods such as altering real data[19] and generating novel synthetic training data[20]. In the latter case, generative AI methods, such as generative adversarial networks[21] and diffusion models[22], could be applied to help researchers to create realistic data for species, environments or scenarios where data are scarce. For example, data could be generated across the categories of interest[23], targeted to specifically fill data gaps (for example, rare species[24]) or could increase the contextual diversity for training or evaluating models in new regions or ecosystems[25]. To this end, generative AI models have already been used to improve the classification of rare mammals in camera traps[26], insects in sticky paper trap images[27] and species in citizen-science data[28]. Moreover, non-generative AI classification models trained with generative-AI-augmented datasets have been used to access new ecological insights from big data, helping us to better understand challenging-to-observe ecological processes. Such models have, for example, been used to identify bat species within vertical-looking radar, which had hitherto been impossible due to their morphological similarities with birds, thereby allowing us to gain new

 https://doi.org/10.1038/s41559-024-02623-1

BOX 1

Underpinnings of generative AI models

Generative AI encompasses a range of data-generation AI models, each adopting unique methods to generate data. For example, in the below figure, we illustrate how generative adversarial networks learn to generate data, such as animal vocalizations, through a feedback loop by using random noise to create fake data that are then evaluated by a discriminating component of the model. The feedback from the discriminator improves the generation of subsequent more realistic-looking data[21]. By contrast, diffusion models learn to generate data, such as camera trap images, by iteratively adding Gaussian noise to a data sample and then gradually denoising the sample over a series of mirroring steps[22]. Note that, often, the underlying architecture used for generating data can be applied to different data types; for example, both generative adversarial networks and diffusion models can be applied to similar data, including audio, images and video. As such, which model architectures to use depends on the specific details of the applications, such as end-use cases and training data quantities, and requires consultation with computer scientists. An overview of how other generative AI models, such as variational autoencoders and large language models, represent distinct approaches to creating new data is included in refs. 80,81.

Often, these models learn the underlying patterns giving rise to real-world data through unsupervised learning, an approach where models are provided unlabelled training data, thus allowing them to take advantage of large datasets where manual human labelling would be too expensive or time consuming. This contrasts with many (though not all) non-generative AI approaches where annotated training datasets are typically required; moreover, generative AI models often require scales of training data several magnitudes greater than non-generative AI approaches. However, when trained with enough data, generative AI algorithms can give rise to foundation models: large-scale models suitable for a range of general tasks, such as image generation across different use cases (for example, ref. 8). Foundation models can further be adapted or fine tuned for improved performance in domain-specific applications with relatively little additional training data. Foundation image-generation models could, for example, be fine tuned to generate camera trap-style images of rare species in undersampled regions (Augmenting data-scarce datasets section), whereas general-purpose large language models could be fine tuned to increase their performance in extracting information from scientific text (Enhancing the accessibility of existing datasets section).

Moreover, different data types (such as text, image or video) can be combined during the training process, and models can learn to make useful connections between diverse data modalities. For example, many image-generation models can combine user-provided text prompts and image inputs to enable customizable image editing or to generate new images[8]. This flexibility across data types is particularly relevant in ecology, where the burgeoning availability of diverse data—from genetic sequences to animal movement tracks—opens the possibility for ecologists to integrate large, diverse datasets within such multimodal generative AI models.

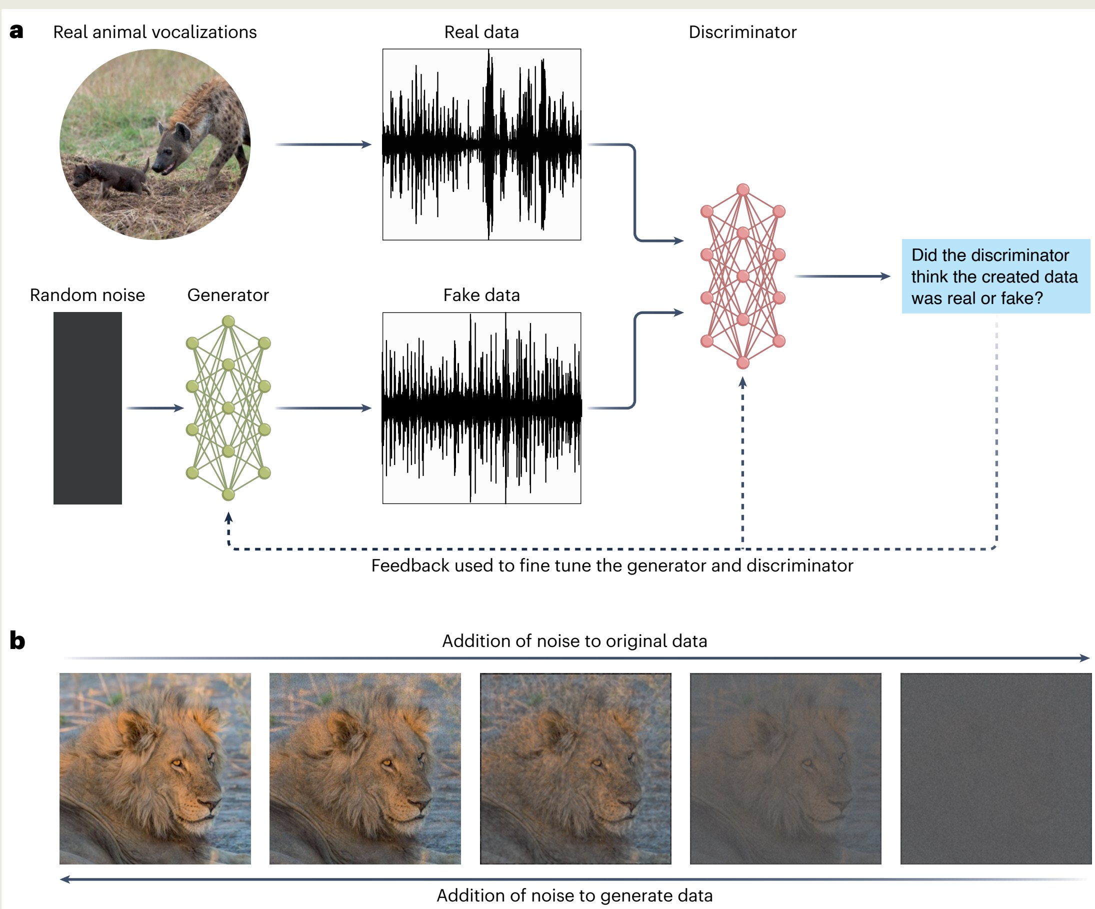

Table 2 | Examples of potential generative AI applications within ecology

Potential application in ecology	Illustrative example	Potential advances to current methods	Parallels in other fields	Key references
Enhancing our understanding of complex ecological processes	Using generative models to understand high-order interactions within plant communities	Could capture complex patterns that are difficult to model heuristically and with other computational approaches	Generative AI is being used within the neurosciences to enhance our understanding of complex brain processes, including brain ageing and disease progression	35,70
Improved modelling of individual dynamics in agent-based models	Using generative agent-based models to study shifts in animal movements and disease risks under global change	Could more accurately represent animal decision-making compared to existing agent-based methods	In the social sciences, generative agent-based models are being used to study complex social behaviours and decisions.	42,43
Forecasting species responses to novel conditions	Using generative AI models to generate probabilistic forecasts of species migrations under extreme climate events	Could enable the prediction of complex behaviours or ecosystem dynamics in new contexts with few examples in the training data	Generative models have been used to predict the pathogenicity of new disease variants and more efficiently quantify uncertainty in climate forecasts	71,72
Simulating population data to infer genetic parameters, such as mutation rates	Using generative models to simulate the effect of habit fragmentation on gene flow	Generative models could perhaps identify genetic parameters from real data, avoiding the biases of setting parameters manually	Generative models have been used to reconstruct the genetic parameters of human data	73
Enriching ecological datasets that have limited data	Augmenting datasets of rare plant species with generated images to improve species classifications in citizen-science applications	Could provide more realistic synthetic data than other methods of data augmentation	Generative models are enhancing medical imaging datasets for rare conditions	74
Annotating and labelling large social–ecological datasets	Using large language models to analyse social media sentiments on attitudes towards nature topics, such as biodiversity loss	Increased efficiency over human screening alone; supports multilingual text; more flexibility in AI text categorization	Generative models are being used by businesses to analyse customer reviews	75
Increasing the scope and scale of literature reviews and meta-analyses	Applying large language models and retrieval-augmented generation to screen papers and extract data for global literature reviews	Increased efficiency over human screening alone; can work with manuscripts from multiple languages	Large language models have been used for screening medical papers	76
Reducing publishing barriers for non-native English writers	Using large language models for writing support and grammar correction in paper submissions	Expands access to writing improvement tools that are currently restricted by scope, location and cost.	Academic publishers are integrating generative AI into their professional journal-editing services	77
Expanding science outreach efforts with the public	Using large language model-powered chatbots to communicate research to underrepresented groups	Allows a greater scope of personalized and interactive science outreach; supports multilingual outreach	AI chatbots have been used to increase engagement with museum visitors	78
Compiling data for policymakers and applied use cases	Using large language models to compile and summarize studies relevant to urban planning	Could improve the ability of non-scientists to find and understand research relevant to their applications; supports multilingual outreach	Large language models are being evaluated in the context of supporting the uptake of nature-based solutions.	79

insights from existing data into large-scale aerial migrations, such as on aerial habitat selection and shifting migration timings[29].

Whereas current ecological data augmentation applications have predominantly focussed on images, these techniques could also be expanded to other data types, such as audio and animal-borne sensor data. For example, Wan and Dodge used generative AI to interpolate data gaps in animal trajectories and found such models performed better than other commonly used interpolation techniques[30], hinting at their potential to recover information from often irregular datasets within movement ecology. Moreover, within engineering, generative adversarial networks have been used to generate synthetic acceleration datasets, enhancing the ability of non-generative AI models to detect building damage[31]. Similarly, in animal ecology, a compelling question is whether generative AI models could augment sparse behaviour occurrences in datasets from animal-borne sensors and thus improve the ability of non-generative AI classification models to recognize important animal behaviours, such as predation events, human encounters and other interspecific interactions[32].

The application of generative AI for data augmentation is relatively new, and challenges exist. For example, in image-generation, current systems struggle to accurately generate images of rare species from only a few examples[33] and can hallucinate or exaggerate species' textural features or morphology[28]. Such hallucinations will also probably extend to the generation of other data types. Moreover, training with synthetically augmented data can be highly inefficient compared to real data, sometimes requiring 10× to 100× more training images to achieve the same performance gains[24,34]. Continuing collaborations between ecologists and computer scientists will be critical for addressing these challenges in the context of ecological applications and in growing the range of generative AI applications for ecological data augmentation.

Extending ecological patterns beyond observed data

Ecological patterns and processes arise from complex interactions between organisms and their environments across multiple spatial and temporal scales. These interactions can be high-dimensional and nonlinear, such that ecological phenomena are often difficult to accurately predict and observe directly[35]. Deciphering the underlying drivers behind these ecological patterns is a central challenge for ecologists due in part to the innate complexities of natural systems, the volumes of data—from animal movements to physiological states to satellite imagery—required to quantify ecological processes and the tractability of contemporary models[36,37].

https://doi.org/10.1038/s41559-024-02623-1

In the face of these challenges, a central question is whether and how generative AI could be more broadly applied to ecological data to help to understand complex interactions between organisms and their environments. One emerging possibility is that generative AI models could help to elucidate ecological relationships by generating new data that reflect and extend real-world patterns present in their training datasets, though this out-of-sample generation is an emerging area of research. For example, whereas generative adversarial networks have been used to generate realistic movement trajectories of central place foragers[38], the inclusion of environmental covariates in the generative process remains a crucial next step for potentially generating trajectories across novel environmental conditions. Additionally, although generative AI has demonstrated an ability to generate outputs that extend beyond training data in other domains, such as language and image generation (for example, ref. 21), similar out-of-sample data generation remains to be empirically validated for many ecological data types (but refer to ref. 35).

The application of generative AI in this context could resemble the role of simulations within ecological modelling, wherein complex simulations can be used to compare against real-world systems for prediction and hypothesis testing[39]. Traditional statistical and mechanistic models, such as those historically used for predicting population growth or resource depletion, are powerful tools often based on strong theoretical foundations and empirical data; generative AI models differ in their method of creating new data. For example, whereas mechanistic approaches often rely on predefined equations to model ecological processes and make predictions, generative AI approaches use AI algorithms optimized for data generation to self-learn directly from the training data the underlying patterns that give rise to real-world observations (Box 1). Such models could, therefore, potentially capture and replicate complex, statistically identifiable processes that would be challenging for ecologists to formally define and parameterize within simulations. As such, generative AI could be used to generate more nuanced representations of ecological phenomena under a broader range of conditions, thereby providing a complementary tool to traditional methods. However, it is important to note that the effectiveness of generative AI in generating realistic data is heavily dependent on the availability of a large and varied training dataset, a limitation that traditional statistical and mechanistic models do not face as acutely. Moreover, caution is necessary when interpreting model outputs, particularly as the risk of generating data that incorrectly represents ecological phenomena is heightened by the opaque nature of the AI data-generation process. Ultimately, this remains an area of emerging research and one where we speculate collaborations between ecologists and computer scientists could lead to exciting discipline-specific applications. For example, models trained with animal movement and environmental data could be used to create forecasts of animal movements across a range of hypothetical climate scenarios, infer ecological connectivity in data-scarce regions and stress-test conservation management strategies.

Furthermore, by analysing the outputs from generative AI models, researchers could develop and test hypotheses about real-world ecological interactions and mechanisms, leading to a deeper understanding of the underlying processes driving ecological patterns. For instance, Hirn et al. used generative AI models to create novel patches of plant species communities that had similar properties to real communities and used the resulting conditional probabilities of species co-occurrences to infer complex indirect interactions involving multiple species[35]. The authors concluded that including data on hypothesized mediators of species interactions within training datasets—such as phenotypic traits—could allow for hypothesized mechanisms driving species co-occurrence to be tested. This could be achieved by generating species compositions across conditions of the mediator that may be hard to empirically observe and statistically testing how these conditions influence the likelihood of different species co-occurring, thereby gaining insights into complex processes driving such ecological dynamics (more details in refs. 35,40).

Integrating generative AI with mechanistic models

Additional potential lies in integrating generative AI directly within existing ecological analytical approaches. For example, agent-based models, which simulate interactions of individual 'agents' with their abiotic and biotic environments, are a foundational tool within ecology used to test diverse ecological theories ranging from migration timing to parental investment[41]. Extensions of these models that integrate generative AI with mechanistic modelling are being used in the social sciences to mimic complex human behaviours more accurately and consequently model their consequences on social dynamics, such as epidemic spread[42–44]. Combined with the growing scale of animal behaviour data available, including from wearable sensors and drones, similar generative agent-based model approaches within ecology could provide exciting opportunities for a more realistic representation of animal decision-making. Such models could enable the testing of complex theories in behavioural ecology, such as the role of personal memory in territory formation and the role of individual decisions in driving emergent group structures within social species[45,46]. For example, analogous to efforts in the social sciences, coupling a mechanistic model of disease transmission with a generative AI model simulating realistic animal behaviours across scenarios could provide deeper insights into the complex, nonlinear interactions that drive disease spread. For examples of how similar methods are being considered across fields, we recommend refs. 43,44.

Ultimately, the potential for generative AI to generate realistic ecological, environmental and behavioural data under novel conditions provides intriguing opportunities for predicting ecosystem dynamics under global change. For instance, within the domain of disease ecology, generative AI could be used to forecast the emergence and transmission of zoonotic diseases due to changes in species distributions and expansion of wildlife–urban interfaces[47], whereas in the context of human–wildlife conflict, generative models could lead to improved predictions of conflict hotspots under shifts in anthropogenic land use and climate change[48]. Generative AI could thus serve as a versatile tool for ecologists, offering a platform for theoretical exploration, hypothesis generation and modelling. These findings, in turn, could inform empirical research and ecosystem management strategies, such as guiding restoration efforts, predicting areas of future conservation concern and aiding invasive species management.

Enhancing the accessibility of existing datasets

In addition to directly modelling ecological phenomena, generative AI could also help to identify key environmental processes underlying the relationships between landscapes and the species they support by increasing the accessibility of environmental data. For example, generative AI models are increasingly being used within the environmental sciences to identify and extract geographic attributes—such as vegetation cover, water bodies and anthropogenic changes—from remote sensing data and can do this with fewer training samples than required for traditional AI techniques[49,50]. Moreover, generative AI models have been used to reduce the dimensionality of complex environmental datasets without the linear dimension reduction constraints often imposed by methods such as principal component analysis, for example, to simplify the modelling of species distributions, thereby streamlining the acquisition of information for conservation planning[51]. In such applications, evaluation of model outputs remains critical to understand whether and where generative AI models offer improvements over less complex approaches.

Additionally, large language models and retrieval-augmented generation, a technique for improving the ability of large language models to retrieve information from specific sources, offer promise in allowing researchers to more efficiently query and extract ecological data from written datasets[52,53]. Such applications could also hold promise for tackling long-standing barriers to diversity, equity and inclusion within ecology, for example, by facilitating the extraction of information from non-English texts, thereby increasing the visibility of underrepresented groups in ecology. For further reading on the potential implications of generative AI in diversity, equity and inclusion, we recommend refs. 54,55 as starting resources.

Limitations of generative AI

Although generative AI presents promising prospects for ecology, it also carries limitations, challenges and risks (such as those touched on throughout) that must be carefully weighed by researchers before its use. Below, we briefly discuss three particularly relevant considerations for ecologists.

Model biases

One important concern is the potential for AI-generated outputs to mirror biases within their training data[56,57]. Biases within generative AI models are systematic errors or misrepresentations that favour certain groups or ideas, perpetuate stereotypes or lead to incorrect assumptions[56]. This issue is particularly relevant for ecology due to geographic, taxonomic and social biases within ecological datasets[2,58–61]. For example, in the context of large language models, it is crucial to guard against perpetuating existing biases against underrepresented groups of researchers, particularly in how models might prioritize or value studies. This could manifest in the unfair ranking or selection of research literature based on their similarities with the predominantly gendered global north bias in ecology publications[59,61] and could lead to models overlooking critical insights from researchers from underrepresented regions or groups. Moreover, issues can arise from the model's creation of fictitious content, which in the context of scientific writing could include fictitious literature citations[62,63]. As generative AI evolves, improvements in bias and reductions in the occurrence of such hallucinations are anticipated[62]. To further mitigate these risks, it is essential to use diverse datasets during model training, ensure active researcher involvement in the review of model outputs and rigorously assess outputs for bias.

Environmental effects

The environmental footprint of AI technologies, both from operational energy consumption and hardware manufacturing[64], poses a dilemma for ecologists committed to sustainability. Generative AI models are particularly energy intensive, with substantial effects on carbon emissions[65,66]. This underscores the need for transparency when integrating AI into ecology and the development of efficiency-focussed algorithms optimized to reduce carbon emissions and energy use[2,65]. Such innovations in AI efficiency also tend to reduce computational complexity and will thus make a broader range of algorithms available for ecologists to run on their own devices without specialist hardware, increasing the accessibility of the technology as a whole.

Privacy and ownership

Generative AI also brings unique data privacy and ownership concerns[67,68], which are particularly pertinent within ecology as data collection is often challenging and resource intensive. Moreover, generative AI models can reproduce specific contents of their training data in their outputs[69], and thus, researchers must be mindful of privacy and security risks. This concern is particularly acute in cases where models are trained with personal data, for example, within social–ecological studies, where there is a risk of models outputting sensitive information[69]. Such privacy concerns also extend to issues of species safety, where the unintended disclosure of sensitive information, such as threatened species locations, could put species at risk.

Concluding remarks

Generative AI is rapidly gaining traction across academic domains and holds immense potential in accelerating the field of ecology. To date, these models have shown preliminary promise in augmenting data-scarce datasets, understanding ecological processes and streamlining the retrieval of complex information. However, the rapidly evolving nature of generative AI means that the scope of its applications within ecology is expanding greatly, and there remains uncertainty on its future role within the field. Whereas initial use cases have focussed on data types such as images and human language, there are emerging opportunities to apply generative AI models to a multitude of ecological data, including, but not limited to, genetic information, animal-borne sensor data and remotely sensed environmental data. However, we acknowledge that with these opportunities come risks and challenges. Issues such as inherent biases in training data, the environmental footprint of running complex AI models and ethical concerns around data privacy and ownership must be carefully navigated. The future of generative AI in ecology lies in the development of robust interdisciplinary collaborations between ecologists and computer scientists. Such partnerships will be important for embedding ecological knowledge within AI and AI within the ecological research process. This will lead to more ecologically meaningful and relevant models and will be critical for leveraging the power of generative AI to drive ecological insights into species across the globe.

References

1. Christin, S., Hervet, É. & Lecomte, N. Applications for deep learning in ecology. *Methods Ecol. Evol.* **10**, 1632–1644 (2019).
2. Tuia, D. et al. Perspectives in machine learning for wildlife conservation. *Nat. Commun.* **13**, 792 (2022).
3. Norouzzadeh, M. S. et al. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proc. Natl Acad. Sci. USA* **115**, E5716–E5725 (2018).
4. Pichler, M., Boreux, V., Klein, A.-M., Schleuning, M. & Hartig, F. Machine learning algorithms to infer trait-matching and predict species interactions in ecological networks. *Methods Ecol. Evol.* **11**, 281–293 (2020).
5. Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
6. Walraven, E., Spaan, M. T. J. & Bakker, B. Traffic flow optimization: a reinforcement learning approach. *Eng. Appl. Artif. Intell.* **52**, 203–212 (2016).
7. Shahriar, S. et al. Putting GPT-4o to the sword: a comprehensive evaluation of language, vision, speech, and multimodal proficiency. *Appl. Sci.* **14**, 7782 (2024).
8. Tian, Y. et al. StableRep: synthetic images from text-to-image models make strong visual representation learners. In *Proc. 37th International Conference on Neural Information Processing Systems* 48382–48402 (Curran Associates, 2024).
9. Esser, P. et al. Structure and content-guided video synthesis with diffusion models in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* 7312–7322 (IEEE, 2023).
10. Naveed, H. et al. A comprehensive overview of large language models. Preprint at *arXiv* https://doi.org/10.48550/arXiv.2307.06435 (2024).
11. Lam, R. et al. Learning skillful medium-range global weather forecasting. *Science* **382**, 1416–1421 (2023).
12. Lin, Z. et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**, 1123–1130 (2023).
13. Nathan, R. et al. Big-data approaches lead to an increased understanding of the ecology of animal movement. *Science* **375**, eabg1780 (2022).
14. Sugai, L. S. M., Silva, T. S. F., Ribeiro, J. W. Jr & Llusia, D. Terrestrial passive acoustic monitoring: review and perspectives. *BioScience* **69**, 15–25 (2019).

15. Beery, S. in *AI for Social Impact* (eds Tambe, M. et al.) https://www.ai4sibook.org/wp-content/uploads/2022/08/MegaDetector.pdf (2022).
16. Hart, A. G. et al. Assessing the accuracy of free automated plant identification applications. *People Nat.* **5**, 929–937 (2023).
17. Beery, S., Van Horn, G. & Perona, P. Recognition in terra incognita. In *Proc. European Conference on Computer Vision* (eds. Ferrari, V. et al.) 456–473 (2018).
18. Koh, P. W. et al. WILDS: a benchmark of in-the-wild distribution shifts. In *Proc. 38th International Conference on Machine Learning*, 139 (ICML, 2021).
19. van Dyk, D. A. & Meng, X.-L. The art of data augmentation. *J. Comput. Graph. Stat.* **10**, 1–50 (2001).
20. Hittmeir, M., Ekelhart, A. & Mayer, R. On the utility of synthetic data: an empirical evaluation on machine learning tasks. In *ARES '19: Proceedings of the 14th International Conference on Availability, Reliability and Security*, article no. 29 (2019).
21. Goodfellow, I. et al. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems* 2672–2680 (Curran Associates, 2014).
22. Dhariwal, P. & Nichol, A. Diffusion models beat GANs on image synthesis. In *Proc. 35th International Conference on Neural Information Processing Systems*, https://doi.org/10.5555/3540261.3540933 (Curran Associates, 2021).
23. Azizi, S. et al. Synthetic data from diffusion models improves ImageNet classification. *Trans. Mach. Learn. Res.* (2023).
24. Beery, S. et al. Synthetic examples improve generalization for rare classes. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)* 852–862 (IEEE, 2020).
25. Vendrow, J., Jain, S., Engstrom, L. & Madry, A. Dataset interfaces: diagnosing model failures using controllable counterfactual generation. In *Data-Centric Machine Learning Research (DMLR) Workshop at ICML 2023*, https://dmlr.ai/assets/accepted-papers/49/CameraReady/icml.pdf (2023).
26. Ghiasi, G. et al. Simple copy-paste is a strong data augmentation method for instance segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, https://doi.org/10.1109/CVPR46437.2021.00294 (IEEE, 2021).
27. Lu, C.-Y., Arcega Rustia, D. J. & Lin, T.-T. Generative adversarial network based image augmentation for insect pest classification enhancement. *IFAC Pap.* **52**, 1–5 (2019).
28. Schwartz, I. et al. Discriminative class tokens for text-to-image diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* 22668–22678 (IEEE, 2023).
29. Werber, Y., Sextin, H., Yovel, Y. & Sapir, N. BATScan: a radar classification tool reveals large-scale bat migration patterns. *Meth. Ecol. Evol.* **14**, 1764–1779 (2023).
30. Wan, Z. & Dodge, S. A generative trajectory interpolation method for imputing gaps in wildlife movement data. In *GeoWildLife 23 Proc. 1st ACM SIGSPATIAL International Workshop AI-Driven Spatio-Temporal Data Analysis Wildlife Conservation* 1–8 (ACM, 2023).
31. Luleci, F., Catbas, F. N. & Avci, O. Generative adversarial networks for labeled acceleration data augmentation for structural damage detection. *J. Civ. Struct. Health Monit.* **13**, 181–198 (2023).
32. Jiguet, F. et al. Multi-sensor data loggers identify the location and timing in four poaching cases of the endangered Eurasian Curlew *Numenius arquata*. *Forensic Sci. Int. Anim. Environ.* **4**, 100069 (2023).
33. Gal, R. et al. An image is worth one word: personalizing text-to-image generation using textual inversion. In *Proc. 11th International Conference on Learning Representations*, https://openreview.net/forum?id=NAQvF08TcyG (2023).
34. Yuan, J. et al. Real-fake: effective training data synthesis through distribution matching. In *Proc. 12th International Conference on Learning Representations*, https://openreview.net/pdf?id=svIdLLZpsA (2023).
35. Hirn, J. et al. A deep generative artificial intelligence system to predict species coexistence patterns. *Meth. Ecol. Evol.* **13**, 1052–1061 (2022).
36. Borowiec, M. L. et al. Deep learning as a tool for ecology and evolution. *Meth. Ecol. Evol.* **13**, 1640–1660 (2022).
37. Chesson, P. General theory of competitive coexistence in spatially-varying environments. *Theor. Popul. Biol.* **58**, 211–237 (2000).
38. Roy, A., Fablet, R. & Bertrand, S. L. Using generative adversarial networks (GAN) to simulate central-place foraging trajectories. *Meth. Ecol. Evol.* **13**, 1275–1287 (2022).
39. Peck, S. L. Simulation as experiment: a philosophical reassessment for biological modeling. *Trends Ecol. Evol.* **19**, 530–534 (2004).
40. Hirn, J., Sanz, V. & Verdú, M. in *Generative AI to Understand Complex Ecological Interactions* 293–308 (Springer, 2024).
41. DeAngelis, D. L. & Diaz, S. G. Decision-making in agent-based modeling: a current review and future prospectus. *Front. Ecol. Evol.* **6**, 237 (2019).
42. Ghaffarzadegan, N., Majumdar, A., Williams, R. & Hosseinichimeh, N. Generative agent-based modeling: an introduction and tutorial. *Syst. Dyn. Rev.* **40**, e1761 (2024).
43. Williams, R., Hosseinichimeh, N., Majumdar, A. & Ghaffarzadegan, N. Epidemic modeling with generative agents. Preprint at *arXiv* https://doi.org/10.48550/arXiv.2307.04986 (2023).
44. Park, J. S. et al. Generative agents: interactive simulacra of human behavior. In *Proc. 36th Annual ACM Symposium on User Interface Software and Technology* 1–22 (Association for Computing Machinery, 2023).
45. Potts, J. R. & Lewis, M. A. How memory of direct animal interactions can lead to territorial pattern formation. *J. R. Soc. Interface* **13**, 20160059 (2016).
46. Strandburg-Peshkin, A., Farine, D. R., Crofoot, M. C. & Couzin, I. D. Habitat and social factors shape individual decisions and emergent group structure during baboon collective movement. *eLife* **6**, e19505 (2017).
47. Han, B. A. et al. A synergistic future for AI and ecology. *Proc. Natl Acad. Sci. USA* **120**, e2220283120 (2023).
48. Abrahms, B. et al. Climate change as a global amplifier of human–wildlife conflict. *Nat. Clim. Change* **13**, 224–234 (2023).
49. Hennessy, A., Clarke, K. & Lewis, M. Generative adversarial network synthesis of hyperspectral vegetation data. *Remote Sens.* **13**, 2243 (2021).
50. Requena-Mesa, C. et al. in *Pattern Recognition* (eds Fink, G. A. et al.) 203–217 (Springer, 2019).
51. Dinnage, R. How many variables does Wordclim have, really? Generative A.I. unravels the intrinsic dimension of bioclimatic variables. Preprint at *bioRxiv* https://doi.org/10.1101/2023.06.12.544623 (2023).
52. Castro, A., Pinto, J., Reino, L., Pipek, P. & Capinha, C. Large language models overcome the challenges of unstructured text data in ecology. *Ecol. Inform.* **82**, 102742 (2024).
53. Gao, Y. et al. Retrieval-augmented generation for large language models: a survey. Preprint at *arXiv* https://doi.org/10.48550/arXiv.2312.10997 (2024).
54. Cooper, N. et al. Harnessing large language models for coding, teaching and inclusion to empower research in ecology and evolution. *Methods Ecol. Evol.* **15**, 1757–1763 (2024).
55. Suri, S. Defining our future with generative AI. *Nat. Comput. Sci.* **4**, 641–643 (2024).
56. Ferrara, E. Should ChatGPT be biased? Challenges and risks of bias in large language models. *First Monday* **28**, 11 (2023).

57. Schramowski, P., Turan, C., Andersen, N., Rothkopf, C. A. & Kersting, K. Large pre-trained language models contain human-like biases of what is right and wrong to do. *Nat. Mach. Intell.* **4**, 258–268 (2022).
58. Konno, K. et al. Ignoring non-English-language studies may bias ecological meta-analyses. *Ecol. Evol.* **10**, 6373–6384 (2020).
59. Maas, B. et al. Women and global south strikingly underrepresented among top-publishing ecologists. *Conserv. Lett.* **14**, e12797 (2021).
60. Clark, J. A. & May, R. M. Taxonomic bias in conservation research. *Science* **297**, 191–192 (2002).
61. Hughes, A. C. et al. Who is publishing in ecology and evolution? The underrepresentation of women and the Global South. *Front. Environ. Sci.* **11**, 1211211 (2023).
62. Azamfirei, R., Kudchadkar, S. R. & Fackler, J. Large language models and the perils of their hallucinations. *Crit. Care* **27**, 120 (2023).
63. Boyko, J. et al. An interdisciplinary outlook on large language models for scientific research. Preprint at *arXiv* https://doi.org/10.48550/arXiv.2311.04929 (2023).
64. Gupta, U. et al. Chasing carbon: the elusive environmental footprint of computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* 854–867 (IEEE, 2021).
65. Chien, A. A. et al. Reducing the carbon impact of generative AI inference (today and in 2035). In *Proc. 2nd Workshop on Sustainable Computer Systems* 1–7 (ACM, 2023).
66. Crawford, K. Generative AI's environmental costs are soaring — and mostly secret. *Nature* **626**, 693 (2024). –693.
67. Elkin-Koren, N., Hacohen, U., Livni, R. & Moran, S. Can copyright be reduced to privacy? In *5th Symposium on Foundations of Responsible Computing (FORC 2024)* 3:1–3:18 (Schloss Dagstuhl – Leibniz-Zentrum fürInformatik, 2024).
68. Smits, J. & Borghuis, T. in *Law and Artificial Intelligence: Regulating AI and Applying AI in Legal Practice* (eds Custers, B. & Fosch-Villaronga, E.) 323–344 (T.M.C. Asser Press, 2022).
69. Huang, J., Shao, H. & Chang, K. C.-C. Are large pre-trained language models leaking your personal information? In *Findings of the Association for Computational Linguistics* 2038–2047 (Association for Computational Linguistics, 2022).
70. Wang, R. et al. Applications of generative adversarial networks in neuroimaging and clinical neuroscience. *NeuroImage* **269**, 119898 (2023).
71. Frazer, J. et al. Disease variant prediction with deep generative models of evolutionary data. *Nature* **599**, 91–95 (2021).
72. Li, L., Carver, R., Lopez-Gomez, I., Sha, F. & Anderson, J. Generative emulation of weather forecast ensembles with diffusion models. *Sci. Adv.* **10**, eadk4489 (2024).
73. Yelmen, B. et al. Creating artificial human genomes using generative neural networks. *PLoS Genet.* **17**, e1009303 (2021).
74. Chen, Y. et al. Generative adversarial networks in medical image augmentation: a review. *Comput. Biol. Med.* **144**, 105382 (2022).
75. Thukral, V., Latvala, L., Swenson, M. & Horn, J. Customer journey optimisation using large language models: best practices and pitfalls in generative AI. *Appl. Mark. Anal.* **9**, 281–292 (2023).
76. Guo, E. et al. Automated paper screening for clinical reviews using large language models. *J. Med. Internet Res.* https://doi.org/10.2196/48996 (2023).
77. AJE. AJE: English editing and author services for research publishers. *AJE* https://www.aje.com/ (AJE, 2024).
78. Gaia, G., Boiano, S. & Borda, A. in *Museums and Digital Culture: New Perspectives and Research* (eds Giannini, T. & Bowen, J. P.) 309–329 (Springer International, 2019).
79. Richards, D., Worden, D., Song, X. P. & Lavorel, S. Harnessing generative artificial intelligence to support nature-based solutions. *People Nat.* **6**, 882–893 (2024).
80. Bond-Taylor, S., Leach, A., Long, Y. & Willcocks, C. G. Deep generative modelling: a comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 7327–7347 (2022).
81. Gm, H., Gourisaria, M. K., Pandey, M. & Rautaray, S. S. A comprehensive survey and analysis of generative models in machine learning. *Comput. Sci. Rev.* **38**, 100285 (2020).

Acknowledgements

We thank U. Paquet for providing early guidance and feedback on the scope of the manuscript and our institutions for supporting our time in preparing the manuscript. K.R. was supported by the University of Washington eScience Postdoctoral Fellowship and the Washington Research Foundation. Z.H. was funded by the National Science Foundation (grant CCF 2019844).

Author contributions

K.R. conceptualized and led the paper. K.R., B.A., M.S.P. and S.B. contributed to the original draft of the manuscript. K.R., B.A., M.S.P., S.B. and Z.H. provided critical feedback, edits and revisions to the manuscript.

Additional information

Correspondence should be addressed to Kasim Rafiq.

Peer review information *Nature Ecology & Evolution* thanks the anonymous reviewers for their contribution to the peer review of this work.