News & views

G LACZ/IMAGEBROKER/SHUTTERSTOCK

Figure 1 | Zebra finches (*Taeniopygia guttata*). The male is singing.

Animal behaviour

Birds convey complex signals in simple songs

Kate T. Snyder & Nicole Creanza

The quality of a bird's song during courtship can influence whether a male is selected as a mate. An innovative approach using machine learning offers a way to analyse the characteristics of birdsong. **See p.117**

How do you measure a song? Researchers have long attempted to quantify songs to understand bird evolution and behaviour. However, basic ways of measuring these diverse sounds, such as through pitch, duration and loudness, cannot fully capture the songs' intricacy, let alone reveal what properties birds find appealing. On page 117, Alam *et al.*[1] propose a holistic method of measuring song using innovative computational tools. Furthermore, the authors present behavioural evidence that suggests their metric can capture variation that is meaningful to the intended listeners of songs — other birds.

Broadly, male songbirds use their songs to attract a mate. Researchers often quantify variations in songs by counting the unique sounds or measuring changes in pitch, but female birds might have a set of preferences that don't perfectly correspond to differences in any single feature of a song. What researchers have discovered so far is that female preferences differ between species in ways that might favour songs that have certain characteristics[2–7], such as being highly complex, consistently produced, accurately learnt or sung quickly.

By harnessing the potential of machine learning and large-scale song data sets, Alam and colleagues measured song complexity by assessing the "spectral distinctiveness of song syllables" — how different the individual sounds (syllables) in a bird's song are from one another. The authors studied captive zebra finches (*Taeniopygia guttata*), a species in which male birds learn a single song early in life (Fig. 1). They recorded thousands of song repetitions for 49 zebra finches and processed the data using a deep-learning syllable-detection algorithm. Each syllable was then represented as an image of its sound frequencies over time. The authors then used a data-visualization technique[8] called UMAP on the resulting 472,475 syllable images. This algorithm performs dimensionality reduction, distilling the complex information from the images to place each syllable on a 2D map of sounds (Fig. 2), in which repeats of the same syllable are clustered closely together and syllables that are more distinct are generally placed farther apart.

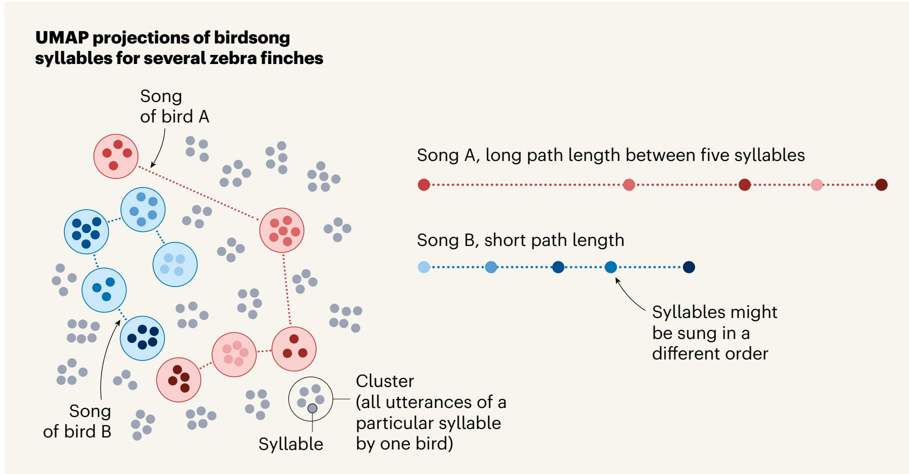

Figure 2 | Analysing birdsong properties. Alam *et al.*[1] present a holistic approach to assess birdsong features in male zebra finches (*Taeniopygia guttata*). Each bird has a single song, which is composed of a varied number of repeats with a core sequence of between two to seven sounds (syllables). The authors used a computational technique called UMAP to plot individual syllables onto a 2D map and found the shortest path between the syllables of a bird's song (the order in which syllables are sung might differ from this). They generated five-syllable songs and females preferred those songs in which the syllables were separated by a longer path length over those with a shorter path length. Songs with a longer path length might be higher-quality songs, although which aspect of the song's qualities the females prefer is not fully understood.

On the resulting UMAP visualization, each unique syllable from each bird formed a discrete cluster that contained the many renditions of that syllable produced over several days of recording. The researchers then 'connected the dots' between each cluster to measure the spectral distinctiveness of a single bird's song as the shortest 'path length' between its different syllable clusters. A long path length means that the bird's song spans more of the possible syllable space, potentially because its syllables are more distinct than are those of songs with a short path length.

To test whether the differences in song captured by the path-length metric are meaningful to female birds, the authors synthesized pairs of five-syllable songs that were different in path length but similar in other acoustic metrics. When these songs were played through separate speakers, females spent more time near the speaker playing the song with the longer path length. This seems to indicate that songs with syllables that are more spread out in the UMAP visualization sound better to female zebra finches.

When birds with songs of different path lengths became fathers and tutored their sons, the juveniles learning short-path-length songs developed songs with a path length similar to or longer than that of their father's songs, whereas those learning long-path-length songs fell short of their father's path length. This finding suggests that long-path-length songs are challenging to learn and thus that path length might be a reliable indicator of a high-quality mate. Together, these outcomes suggest that UMAP path length quantifies biologically meaningful differences between songs in a holistic way, potentially bringing us closer to measuring song as the complex signal it is.

Notably, the distances between syllable clusters do not seem to strongly correlate with differences in conventional acoustic or similarity measurements. Although some dissociation from the raw data is expected with UMAP because of its limitations in preserving the overall data structure[8], this complicates interpretation of the path-length metric. Furthermore, some juveniles in this study learnt their father's song accurately yet still had a shorter path length than their tutor. It is tempting to think that a song with a long path length showcases a bird's virtuoso performance of many diverse song components. However, it remains unclear what it means, acoustically and behaviourally, for two syllables to be far apart on the UMAP visualizations. What are females detecting when they prefer a long-path-length song, and which aspects of these songs do juveniles struggle to learn? How do we reconcile conventional acoustic measurements with these new representations?

These questions raise broader issues about dimensionality-reduction techniques, such as UMAP, which use complex non-linear transformations to project data into a 'space'. Some aspects of UMAP projections can be misleading, especially when representing biological data as being separated into distinct clusters when researchers are, in reality, sampling a continuum. This distinction is highly relevant in human population genetics[9] (see *Nature* https://doi.org/mk6w; 2024) and could also be important for birdsong – for example, when juvenile birds produce highly accurate imitations of their tutor's syllables. Furthermore, one reason there is much scepticism regarding the usefulness of between-cluster distances in UMAP space is because these distances vary when the algorithm is run several times, even when users don't change the parameter settings[10] (see go.nature.com/4cdn2b5). These distances might be more robust if they were produced by averaging path lengths across many UMAP projections, and it would be useful to see whether averaged path lengths yield consistent trends regarding female preference and juvenile learning behaviours.

Machine learning is already revolutionizing the way in which both scientists and the public interact with birdsong, particularly through algorithms for species identification[11] that can be used though mobile-phone apps (see go.nature.com/3pa7vay). However, the user knows a bird-identification algorithm has performed poorly when it predicts an incorrect species.

For a machine-learning algorithm to predict a more nuanced trend, such as whether females will prefer a song, we should be careful to maintain a direct connection between the output of the algorithm and identifiable features of the underlying data. Ideally, other researchers should be able to test the capabilities and limits of any given methodology by retracing the steps from original recordings through the complete coding pipeline and then link the algorithm outputs back to the characteristics that birds actually perceive. Without this connection, we risk our behavioural analyses becoming obscured in a methodological black box that could further confound our understanding of the complex links between animal perception, cognition and behaviour. Nevertheless, developing computational techniques to analyse and compare songs directly, rather than by examining a set of simple acoustic metrics, has the potential to considerably advance the field of behavioural biology and deepen our understanding of animal communication.

Kate T. Snyder and **Nicole Creanza** are in the Department of Biological Sciences, Vanderbilt University, Nashville, Tennessee 37240, USA, and in the Evolutionary Studies Initiative, Vanderbilt University.

e-mails: katherine.t.snyder@vanderbilt.edu; nicole.creanza@vanderbilt.edu

1. Alam, D., Zia, F. & Roberts, T. F. *Nature* **628**, 117–121 (2024).
2. Nowicki, S., Searcy, W. A. & Peters, S. *Proc. R. Soc. Lond. B* **269**, 1949–1954 (2002).
3. Byers, B. E. & Kroodsma, D. E. *Anim. Behav.* **77**, 13–22 (2009).
4. Soma, M. & Garamszegi, L. Z. *Behav. Ecol.* **22**, 363–371 (2011).
5. Robinson, C. M. & Creanza, N. *Ecol. Evol.* **9**, 8362–8377 (2019).
6. Ballentine, B., Hyman, J. & Nowicki, S. *Behav. Ecol.* **15**, 163–168 (2004).
7. Woolley, S. C. & Doupe, A. J. *PLoS Biol.* **6**, e62 (2008).
8. McInnes, L., Healy, J. & Melville, J. Preprint at arXiv https://doi.org/10.48550/arXiv.1802.03426 (2018).
9. National Academies of Sciences, Engineering, and Medicine *et al. Using Population Descriptors in Genetics and Genomics Research: A New Framework for an Evolving Field* (National Academies Press, 2023).
10. Chari, T. & Pachter, L. *PLoS Comput. Biol.* **19**, e1011288 (2023).
11. Kahl, S., Wood, C. M., Eibl, M. & Klinck, H. *Ecol. Inform.* **61**, 101236 (2021).

The authors declare no competing interests.
This article was published online on 20 March 2024.

Engineering

The electronic patch that's impervious to sweat

Yifan Rao & Nanshu Lu

A smart adhesive patch that wicks sweat away from electronics embedded in its centre offers comfortable and reliable sensing of the wearer's biometrics or environment without the risk of perspiration damaging the devices. **See p.84**

Human skin is a remarkable organ that can regulate body temperature by secreting sweat that then evaporates. Sweat-wicking fabrics can draw trapped moisture away from the skin, preventing irritation, but it is complicated to incorporate such materials into wearable electronic patches that contain sensors and circuits for biometric or environmental monitoring. These smart patches must be small and soft, but they also need to resist sweat, which can make the patch peel off and lead to signal loss. Sweat can also seep into the electronics, causing short circuits and corrosion. Although the softness of these patches continues to improve, as does their ability to conform to the wearer's skin[1,2], the challenges of perspiration persist. On page 84, Zhang *et al.*[3] report a flexible device that can discharge sweat rapidly and strategically, ensuring comfort and signal stability during prolonged wear.

Previous designs for wearable electronic patches involve mesh-like electrodes that don't have backing layers, making them very breathable[4,5]. However, the absence of a substrate supporting the electrodes makes it difficult to attach these devices to the skin and to integrate them with advanced electronics. Other designs have incorporated permeable substrates that discharge vapour and sweat through micrometre- or nanometre-sized pores[6–9], but such openings can allow sweat to flow both ways.

These problems can be solved by using fabrics whose wettability changes gradually across their thickness to support the mesh electrodes. Such fabrics facilitate unidirectional liquid transport, because the liquid is driven from the side with a lower surface energy (making it hydrophobic) to the side with a higher surface energy (hydrophilic), whereas flowing in the reverse direction requires too much energy[10]. This concept inspired Zhang *et al.* to engineer the first component of their patch — a device called a vertical liquid diode, which allows liquid to seep from the bottom to the top, but not the other way (Fig. 1).

The authors began their fabrication process by immersing a polyester fabric in a suspension that made it superhydrophobic, and then dip-coated one side of the fabric with a water-based adhesive. The superhydrophobicity of the fabric prevented its porous structure from being compromised by the adhesive. Zhang *et al.* then used a stencil to selectively treat some regions of the other side of the fabric with oxygen plasma (ionized gas), which resulted in these areas becoming hydrophilic.

In this way, the authors created a wettability gradient that spontaneously directed sweat produced at the skin–fabric interface to the top of the fabric in droplet form, effectively preventing any backflow. When optimized, the vertical liquid diode could transport sweat at a rate of 11.6 millilitres per square centimetre per minute, and it remained effective for more than one month. This transport rate is particularly impressive because it is at least 4,000 times faster than the rate at which human sweat is discharged during mild exercise.

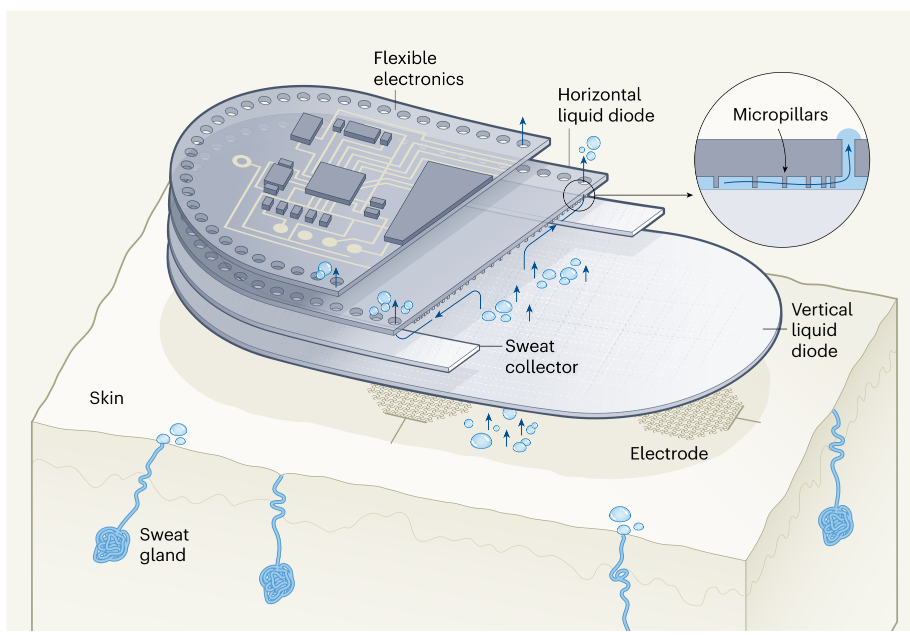

Figure 1 | Sweat-wicking wearable electronics. Zhang *et al.*[3] designed a wearable electronic patch by combining two liquid diodes — devices that allow liquid to flow in one direction only. The vertical liquid diode was made of fabric that was hydrophilic on the top and hydrophobic on the bottom, making sweat produced on the surface of the skin move through the fabric from the bottom to the top. A horizontal liquid diode on top of the vertical one used micropillars distributed in a non-uniform manner to drive the sweat that reached the top of the vertical liquid diode to the edges of the patch, where it was released from outlets. A ring of polyester fabric between the two layers acted as a sweat collector, and all three components were positioned between mesh-like electrodes and a layer of flexible electronics. This smart design is more comfortable than existing wearable electronics, and stabilizes signals produced by the device. (Adapted from Fig. 2a of ref. 3.)