



HAL
open science

Using generative artificial intelligence to test hypotheses about animal signal evolution: A case study in an ornamented fish

Yseult Héjja-Brichard, Kara Million, Julien P Renoult, Tamra C Mendelson

► To cite this version:

Yseult Héjja-Brichard, Kara Million, Julien P Renoult, Tamra C Mendelson. Using generative artificial intelligence to test hypotheses about animal signal evolution: A case study in an ornamented fish. 2023. hal-04276927

HAL Id: hal-04276927

<https://hal.science/hal-04276927>

Preprint submitted on 9 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Using generative artificial intelligence to test hypotheses about animal signal evolution: A case study in an ornamented fish

Yseult Héjja-Brichard¹, Kara Million¹, Julien P Renoult^{2*}, Tamra C Mendelson^{1*}

¹ Department of Biological Sciences, University of Maryland, Baltimore County, Baltimore, Maryland, United States; ² CEFE, Univ Montpellier, CNRS, EPHE, Univ Paul-Valéry Montpellier, Montpellier, Occitanie, France * denotes equal contribution

Abstract

The sensory drive hypothesis of animal signal evolution suggests that animal communication signals evolve in response to environmental pressures. While classical approaches to testing this hypothesis focus on one aspect of the signal, deep learning techniques like generative models can create and manipulate stimuli without targeting a specific feature. Here, we used a technique called style transfer to experimentally test preferences for colour patterns in a fish. We manipulated how similar or dissimilar male body patterns were to their habitats using the Neural Style Transfer (NST) algorithm. We predicted that males whose body patterns are similar to their habitats are easier to process and thus preferred by conspecifics. Our findings suggest that both males and females tend to be sensitive to habitat congruence in their preferences, but to different extents, requiring additional investigation. Nonetheless, this study demonstrates the potential of deep learning techniques in testing hypotheses about animal communication signals.

Introduction

Animals are characterized by a vast array of colours and patterns used for signalling. In the context of mating in particular, understanding the drivers of such diversity in sexual signals has been a central focus of evolutionary biology. Several hypotheses have attempted to explain the evolution of those sexual signals, one of the most influential being the sensory drive hypothesis (Endler, 1992; Endler & Basolo, 1998; Seehausen et al., 2008). Based on signal detection theory, sensory drive suggests that the exploitation of some sensory biases, driven by the environment, leads to the development of different signals. Evidence supporting this hypothesis has been reported in many studies and species. A recent example is the case of the colourful dewlap of Anolis lizards. Fleishman, Perez-Martinez and Leal (2022) analysed the dewlap colours of 17 Anolis species and examined the relationship between colour and habitat light intensity of the different species. They found that red dewlaps are more visible in bright environments and white dewlaps in darker environments.

This result complements previous studies (Fitch & Hillis, 1984; Fleishman, 1992) showing that *Anolis* species with red or orange dewlaps are more often found in bright habitats, while species with white or yellow dewlaps tend to live in darker-lit habitats. This is a strong case in favour of sensory drive, suggesting an adaptation of the dewlap colours to the light environment to enhance the detectability of a body part used in courtship.

Expanding research on sensory drive from simple features like colour to the study of patterns, Hulse and colleagues (2020) assessed the link between the body patterns of colourful fish species (“darters”) and the spatial statistics of their respective habitats. Authors here were testing a “processing bias” hypothesis, which posits that the spatial statistics of natural scenes are attractive because they are the easiest to process by the brain (Renoult & Mendelson, 2019). Thus animal breeding patterns that mimic the spatial statistics of their habitats should be selected by potential mates and therefore evolutionarily favoured. Using the Fourier slope, a measure of the distribution of spatial frequencies (ie luminance distribution) in images, they found differences between males and females: the Fourier slope of male body patterns correlated with the slope of the habitat, whereas female patterns did not. Sexual selection could have favoured males whose body patterning imitates the statistics of their environment as those patterns might be preferred by females.

While these results, showing a correlation between colour or pattern and species habitat, are compelling examples of sensory drive, the approaches are significantly limited in that they focus on one aspect of the signal only, the colour or the spatial frequencies. This is a common pitfall in classical approaches to the study of communication signals, where researchers usually select one feature to investigate, in a more or less arbitrary manner and independently of other features, such as the geometry or symmetry of patterns, colours, or the interaction between multiple features.

Limiting the study of signals to one dimension only fails to reveal the importance of complex patterns or multimodal signalling in animal communication and sexual selection. For instance, Wilson and colleagues (2013) theoretically demonstrated how relying on multimodal signaling is less prone to interpretation errors by the receiver. Similarly, in their review, Pérez-Rodríguez and collaborators (2017) highlighted the need to consider animal color patterns as a whole in complement to individual color features to account better for the multi-component aspect of color patterns.

To address this limitation, tools from artificial intelligence and in particular from deep learning represent a valuable complementary approach. Generative models, such as variational autoencoders (VAEs; Kingma & Welling, 2019), generative adversarial networks (GANs, Goodfellow et al., 2014; 2020) or (latent) diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song & Ermon, 2020), can be used to create or manipulate stimuli without the need to target a specific feature. For instance, CycleGAN (Zhu et al., 2020) is a GAN-based algorithm that can be trained to learn the mapping between two sets of images belonging to different classes (e.g. two types of habitats, or a habitat and the organism that lives there) and be used to transform one image into a different category. Diffusion models can take a sentence as input and turn it into an image (see e.g. (Ramesh et al., 2021; Nichol et al., 2022) and were made famous with the release of the DALL-E application by OpenAI. These models can thus be used to synthesize rich and controlled images to test specific hypotheses about signal preferences. Although they have interesting properties and can

generate high-quality images, such models require a large number of images to be trained, extensive computer power, and they have a somewhat complex architecture. The use of APIs such as DALL-E can be used as an alternative but offers little control to the user.

One tool of artificial intelligence that allows a high level of control over stimulus generation is the Neural Style Transfer (NST) algorithm developed by Gatys and colleagues (2015, 2017). Using the architecture of a convolutional neural network (CNN), the NST algorithm extracts the content (global arrangement of an image, such as the objects that constitutes it) and the style (i.e. texture information, local variations of luminance) of two images, respectively. The algorithm creates a third image that has both the content of one image and the style of the other one (figure 1). Importantly, the user can determine how much weight to put on the style or on the content to modulate the output image as desired. Moreover, the algorithm only requires two input images and no extra training from the user end, and there are several accurate implementations that are ready to use and with tutorials.

Here we used NST to study mate preferences in *Etheostoma caeruleum*, a freshwater fish species native to eastern North America. Following previous results showing a correlation between male body patterns and the species' habitat (Hulse et al., 2020), we asked whether this link could be the result of sexual selection. The processing bias hypothesis stipulates that patterns that are easier to process will be preferred by conspecifics (Renoult & Mendelson, 2019). We thus predicted that males whose body patterns are the closest to their habitat's statistics, thus being easier to process, will be preferred over males that are less similar to their habitats. We manipulated male colour patterns to be either more or less similar to species-specific habitats, using the style transfer technique. We tested the preference of male and female darters for images of conspecifics that do and do not exhibit the "style" of their own habitat.

Besides manipulating the fish patterns, we also alternated the video background of the fish animations to determine whether individuals have a preference for the fish modified with the style of the species-specific habitat, irrespective of the habitat in which it is found (i.e. the background video). We compared the amount of time fish spent with the animated fish exhibiting the species-specific habitat (and with a non-species-specific habitat) when shown in a species-specific habitat vs in a non-species-specific habitat. Finally, we explored the influence of style transfer by comparing preferences for modified vs non-modified animated fish. Our results shed light on the hypothesis that sexual signal design is driven by pattern preferences shaped by the spatial statistics of a species' environment.

Methods

Fish collection and maintenance

We collected males and females of *E. caeruleum* from Israel Creek (39°29'35.1492' N, 77°19'28.0488" W) and Tuscarora Creek (39°15'6.948" N, 77°28'48.288" W) in Frederick County, Maryland between March 15th and April 21st 2022. Fish were transported to the University of Maryland Baltimore County in aerated coolers and individually housed in a recirculating aquarium system (Aquanearing, Inc.). They were provided natural lighting matching their photoperiod with a constant temperature of 19°C and fed a diet of frozen

bloodworms every other day. Fish were euthanized humanely at the end of the study using tricaine methanesulfonate (MS-222).

Permission to collect fish was granted by the Maryland Department of Natural Resources Fishing and Boating Services Scientific Collection (permit number SCP202246A). Housing and experimental procedures were approved by the Institutional Animal Care and Use Committee of the University of Maryland, Baltimore County (protocol 623).

All tested individuals were sexually receptive based on the presence of nuptial colouration in males and visible gravidity in females.

Stimuli creation

We used Neural Style Transfer (NST), a CNN-based algorithm developed by Gatys and colleagues (2017), to alter the spatial statistics of photographs of male *E. caeruleum*. The algorithm requires two input images, one for content and one for style, and aims to render the content image in the style of the style image. The style information is extracted via a Gram matrix while the content information corresponds to the fourth feature layer of a 5-feature layer CNN (figure 1).

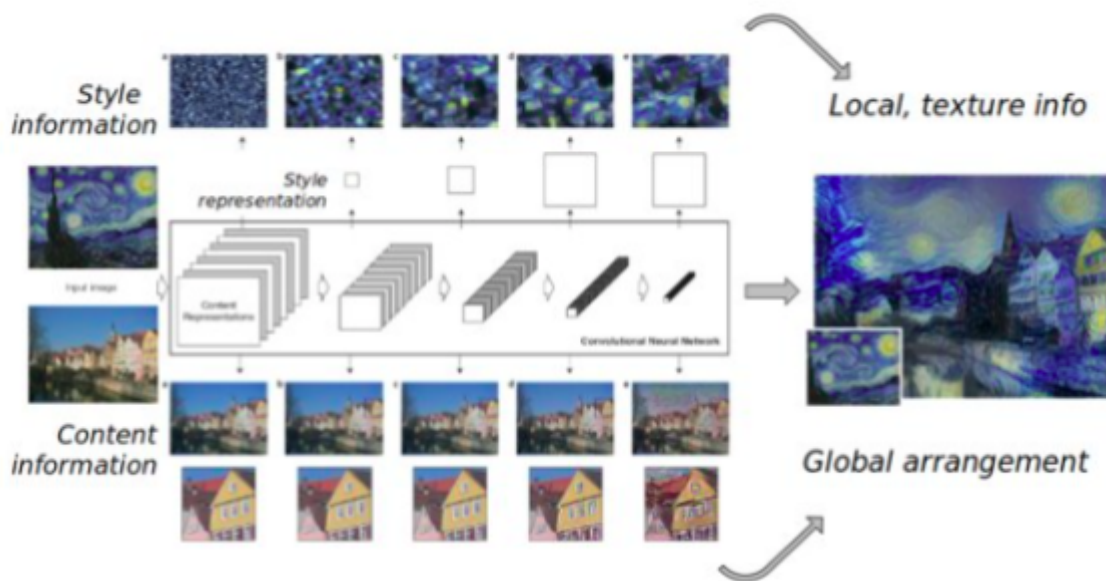


Figure 1. Neural Style Transfer. The algorithm produces a new image by minimizing the distance with both content and style images. The style information corresponds to local statistics of the image (texture, color, paint strokes), while the content information is about the global arrangement of an image (what the image is made of, e.g. the row of houses along the river). Adapted from Gatys et al., 2016.

For our study purposes, a picture of an original fish and a picture of the habitat (either gravel or sand) were fed to the NST algorithm as content and style images, respectively. Content and style have adaptable weights to improve the final result. We kept them constant across our stimuli, using the recommended values (content weight = 1e0; style weight = 1e3). For the style, we only considered the two first layers (out of 5) as they tend to represent the local statistics of the given image.

The original colours of the fish were put back on the final image, using the colour control option of the NST, to make it look more naturalistic. An example is shown in figure 2. Stylized fish were then animated with the Blender software (Blender Online Community, 2018) and displayed in videos depicting either sand or gravel habitats.

The PyTorch (Paszke et al., 2019) implementation of the algorithm (https://pytorch.org/tutorials/advanced/neural_style_tutorial.html) with our modifications as well as the detailed procedure used to create our stimuli are available at: https://github.com/yseulthb/NST_fish.

To control for potential biases, we used three images per habitat category (gravel or sand) and styled five different male individuals. Animated fish were presented against a background video of either gravel (species-specific habitat) or sand (foreign habitat), taken using an Ikelite 200DL underwater housing for a Canon EOS 5D Mark IV digital camera equipped with a Sigma 24 mm f/1.4 lens (see Hulse et al., 2020 for details). We selected three different background videos per habitat. For each pair of video animations presented to test individuals, we matched their luminance values (both in terms of background and animated fish). Pairs of videos also presented different animated individuals. Beyond that constraint, all possible combinations were randomized across trials and across tested individuals.

In short, we applied the style of a species-specific habitat to the body of a male fish in order to present live fish with a choice between animated fish styled in different habitats. We used gravel as the preferred habitat and sand as the foreign habitat as *E. caeruleum* is not typically found on that type of substrate. Conveniently, sand habitats are also the most different from gravel habitats when comparing the spatial frequencies of freshwater habitat types (Hulse et al., 2020).

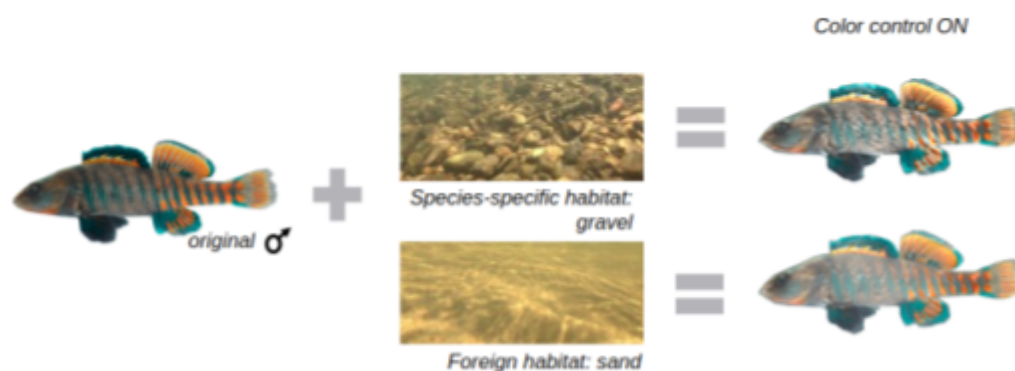


Figure 2: Creation of the stimuli. A picture of an original fish (male individual) and a picture of the habitat (either gravel or sand) are fed to the Neural Style Transfer algorithm as content and style images, respectively. The algorithm extracts the style information via a Gram matrix while the content information corresponds to the usual feature layers of a CNN. The original colours of the fish were put back on the final image to make it look more naturalistic.

Dichotomous choice experiments

We ran a series of six experiments to determine whether *E. caeruleum* showed a preference for individuals modified to exhibit their habitat statistics. Focal fish were presented with a pair of stimuli with different animated fish but the same video background. The different experimental conditions, their characteristics and sample sizes are summarised in Table 1. In addition to estimating preference for differently styled fish, we estimated in a seventh experiment their preference for their own habitat irrespective of the animated fish (ie with the same animated, unstyled fish shown on both sides).

Trials took place in a 37.9L tank flanked by two monitors along the short sides showing the video animations. Water was systematically changed between trials and tanks were filled up to 75% of their capacity. Trials were video-recorded using an overhead video camera (Logitech). A 12cm wide zone in the middle of the tank was delineated with transparent removable dividers and served as an acclimation area. Association zones were defined as the 5cm-wide area in front of the video animations (see Figure 3). This distance corresponds to the usual courting distance between males and females (darters are pair spawners) and is in the range of an estimated visual acuity for that species (Caves et al., 2017).

To control for experimental side bias, each trial alternated on which side of the tank the stimulus conditions (eg gravel-styled or sand-styled) were shown. The tank's lateral (long) sides were blocked with black matte plexiglass to isolate the fish from its surroundings and prevent fish reflection to avoid any distractions. Most fish took part in all seven trials and the experiment order was randomised between individuals to control for potential order effects. Furthermore, the use of 5 different individuals as animated fish combined with the different modification styles limits any recognition bias that could influence our results.

For the first five minutes of a trial, the video animations were blocked with opaque dividers. Those opaque dividers were then removed while the fish was still contained in the acclimation area for an additional five minutes. This gave the focal fish time to visually explore both stimulus sides. Finally, transparent dividers were removed and the focal fish could freely move in the tank for the next 15 minutes. We defined a fish's preferred stimulus by comparing the amount of time spent in the two association zones during the first 12 min after the transparent dividers were removed. Time spent with a stimulus in a dichotomous mate choice trial has been shown to be a reliable proxy for preference in freshwater fishes (Aspbury & Basolo, 2002; Brooks & Endler, 2001; Lehtonen & Lindström, 2008).

Trial videos were scored using the BORIS software (Friard & Gamba, 2016) by four persons that were blind to experimental conditions. Interobserver reliability was tested on 6 trials and estimated to be 90% before equally sharing trials across scorers.

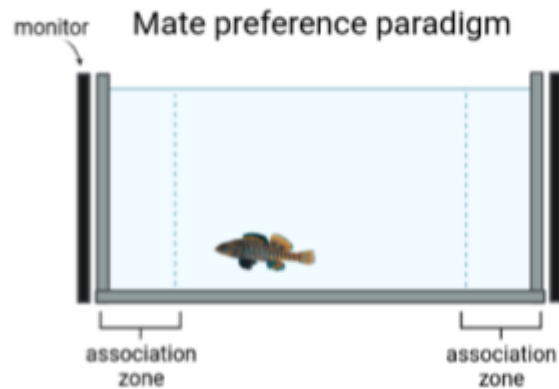


Figure 3: Experimental protocol. The focal fish is shown two different stimuli on monitors situated at both extremities of the tank. The preferred stimulus is defined by comparing time spent in the association zones, the 5cm-wide area in front of the video animations. Figure created with BioRender.com.

Table 1: Experimental conditions and sample sizes for the fish preference experiments

Experiment	Background video	Styled fish 1	Styled fish 2	Sample size
Experiment 1	species-typical	species-typical	foreign	M: 20 F: 17
Experiment 2	foreign	species-typical	foreign	M: 17 F: 16
Experiment 3	species-typical	species-typical	non-modified	M: 19 F: 13
Experiment 4	species-typical	foreign	non-modified	M: 16 F: 15
Experiment 5	foreign	species-typical	non-modified	M: 20 F: 17
Experiment 6	foreign	foreign	non-modified	M: 19 F: 16

Statistical analyses

Statistical analyses were performed with the software R (R Core Team, 2021). For all seven experiments, we first checked for a potential side bias by comparing times spent in the left vs right association zones irrespective of the stimuli shown. We ran either t-tests or Wilcoxon tests depending on the data distribution.

To estimate habitat or fish preference, we compared times spent in the association zone for each condition as well as times spent in the central area. This latter measure informs us of the fish's interest in the stimuli. To compare times spent in those three areas, we ran one-way ANOVAs. If ANOVA results were significant, we ran posthoc tests (Tukey HSD tests) to identify which pairs of stimuli were significantly different. We also report the number of entries in each association zone, a proxy for the fish's mobility. Finally, we counted the number of times fish entered the association zones and used chi-square tests to identify any difference.

All analyses are reported for males and females separately.

Results

Fish preference

We analysed the data of 13 to 17 females and 16 to 20 males (see Table 1) for fish preference. We tested a minimum of 15 individuals per experiment. However, individuals that either froze or stayed in the same location without further moving for the entire duration of the experiment were discarded. We found no side bias for any of the six experiments testing fish preference, neither for females nor males. Detailed test results are provided in the supplementary material.

Experiment 1 tested the preference for fish modified with the statistics of species-typical habitat versus foreign habitat with a background video of species-typical habitat. ANOVA results were significant for females only ($F(2,48) = 4.75$ $p = 0.013$ for females; $F(2,57) = 1.89$ $p = 0.16$ for males). Tukey's HSD Test for multiple comparisons found that the amount of time females spent in the association zones of the fish with species-typical statistics was significantly greater than the amount of time spent in the central zone ($t = -3.03$, $p = 0.011$).

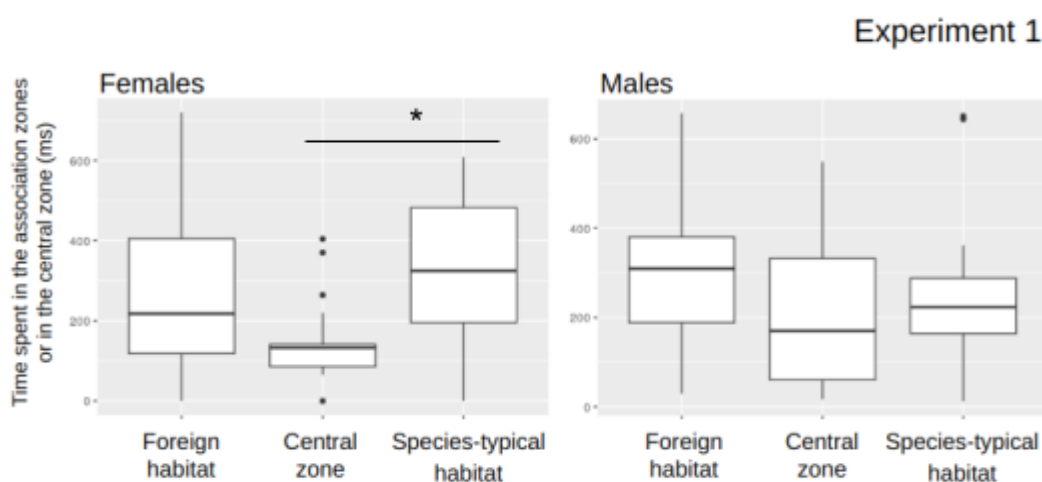


Figure 4: Time (ms) spent in the association zones of the fish with foreign habitat statistics or of the fish modified with species-typical statistics and in the central zones for females (left) and males (right). The video background depicts the species-typical habitat (gravel).

On average, males entered the association zone of the fish modified with species-typical statistics 6.5 times and of the fish with foreign habitat statistics 7.8 times. For females, they entered 6.35 times and 5.53 times, respectively. No significant difference was found for either group ($X^2 = 0.05708$, $df = 1$, $p = 0.811$ for females; $X^2 = 0.1182$, $df = 1$, $p = 0.731$ for males).

Experiment 2 tested a preference for fish modified to reflect the statistics of species-typical habitat versus foreign habitat with a background video of a foreign habitat. Again, ANOVA results were significant for females only ($F(2,45) = 10.3$ $p < 0.001$ for females; $F(2,45) = 10.3$ $p < 0.001$ for males). Tukey's HSD Test for multiple comparisons revealed statistical differences between the time females spent in the association zone of the fish with the statistics of their own habitat ($t = -4.54$, $p < 0.001$) and of the fish modified to reflect a foreign habitat ($t = 2.57$, $p = 0.035$) compared to the central zone.

On average, males entered the association zone of the fish modified with species-typical statistics 7.06 times and of the fish modified to reflect a foreign habitat 8.12 times. For females, it was 4.4 times and 4.2 times, respectively. No significant difference was found for either group ($X^2 = 0.004106$, $df = 1$, $p = 0.949$ for females; $X^2 = 0.07387$, $df = 1$, $p = 0.786$ for males).




Figure 5: Time (ms) spent in the association zones of the fish with foreign habitat statistics or of the fish modified with species-typical statistics and in the central zones for females (left) and males (right). The video background depicts a foreign habitat (sand).

Experiment 3 tested a preference for fish modified to reflect the statistics of species-typical habitat versus a non-modified fish with a background video of the species' habitat. ANOVA results were significant for both males ($F(2,54) = 13.7$, $p < 0.001$) and females ($F(2,36) = 7.94$, $p = 0.0014$). The Tukey test for multiple comparisons found that the amount of time females spent in the association zone of the fish modified to reflect the statistics of their own habitat was significantly different from the amount of time spent in the central zone ($t = -3.96$, $p < 0.001$).

For males, statistical differences were found between times spent in the association zone of the fish modified to reflect the statistics of their own habitat compared to the central zone ($t = -5.17$, $p < 0.001$) and between times spent in the association zone of the fish modified to reflect the statistics of their own habitat compared to the non-modified fish ($t = -3.27$, $p = 0.0052$). Thus **males preferred the fish that was modified to reflect their own habitat over the natural version of the fish.**

On average, males entered the association zone of the fish modified with species-typical statistics 6.3 times and of the fish modified to reflect a foreign habitat 5.2 times. For females, it was 4.9 times and 4.4 times, respectively. No significant difference was found for either group ($X^2 = 0.02308$, $df = 1$, $p = 0.879$ for females; $X^2 = 0.107$, $df = 1$, $p = 0.744$ for males).




Figure 6: Time (ms) spent in the association zones of the non-modified fish or of the fish modified with species-typical statistics and in the central zones for females (left) and males (right). The video background depicts the species-typical habitat (gravel).

Experiment 4 tested a preference for fish modified to reflect the statistics of a foreign habitat versus a non-modified fish with a background video of the species' habitat. ANOVA results were significant for both males ($F(2,45) = 3.45$, $p = 0.04$) and females ($F(2,42) = 6.9$, $p = 0.0026$). The Tukey test for multiple comparisons revealed statistical differences between the time females spent in the association zone of the fish with foreign habitat statistics ($t = 3.65$, $p = 0.0021$) and of the non-modified fish ($t = 2.44$, $p = 0.0487$) compared to the central zone. For males, statistical differences were found between times spent in the association zones of the non-modified fish compared to the central zone ($t = 2.48$, $p = 0.044$).

On average, males entered the association zone of the fish with foreign habitat statistics 6.33 times and of the non-modified fish 6.73 times. For females, it was 6.0 times and 6.06 times, respectively. No significant difference was found for either group ($X^2 = 0.01224$, $df = 1$, $p = 0.912$ for females; $X^2 = 0.0003238$, $df = 1$, $p = 0.986$ for males).




Figure 7: Time (ms) spent in the association zones of the non-modified fish or of the fish with foreign habitat statistics and in the central zones for females (left) and males (right). The video background depicts the species-typical habitat (gravel).

Experiment 5 tested a preference for fish modified to reflect the statistics of a species' habitat versus a non-modified fish with a background video of a foreign habitat. ANOVA results were significant for both males ($F(2,57) = 3.52$ $p = 0.036$) and females ($F(2,48) = 10.2$, $p < 0.001$). Tukey test for multiple comparisons revealed statistical differences between the time females spent with the fish modified to reflect the statistics of their own habitat ($t = -3.86$, $p = 0.00099$) and of the non-modified fish ($t = 3.97$, $p = 0.00073$) compared to the central zone. For males, statistical differences were found between times spent in the association zone of the fish modified to reflect the statistics of their own habitat compared to the central zone ($t = -2.43$, $p = 0.048$).

On average, males entered the association zone of the fish modified with species-typical statistics 7.75 times and of the non-modified fish 8.0 times. For females, it was 4.12 times and 3.94 times, respectively. No significant difference was found for either group ($X^2 = 0.003864$, $df = 1$, $p = 0.95$ for females; $X^2 = 0.003968$, $df = 1$, $p = 0.95$ for males).

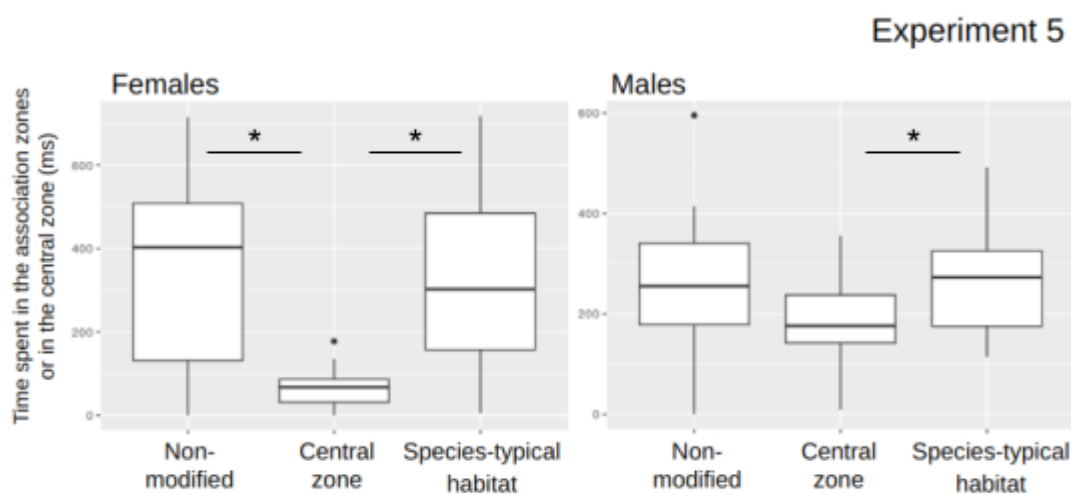


Figure 8: Time (ms) spent in the association zones of the non-modified fish or of the fish modified with species-typical statistics and in the central zones for females (left) and males (right). The video background depicts a foreign habitat (sand).

Experiment 6 tested a preference for fish modified to reflect the statistics of a foreign habitat versus a non-modified fish with a background video of a foreign habitat. ANOVA results were significant for females only ($F(2,45) = 7.24$ $p = 0.0019$ for females; $F(2,54) = 3.01$, $p = 0.058$ for males). Tukey's HSD Test for multiple comparisons found that the amount of time females spent in the association zones of the fish modified to reflect the statistics of a foreign habitat was significantly different from the amount of time spent in the central zone ($t = 3.81$, $p = 0.0013$).

On average, males entered the association zone of the fish with foreign habitat statistics 8.1 times and of the non-modified fish 8.7 times. For females, it was 6.6 times and 5.9 times, respectively. No significant difference was found for either group ($X^2 = 0.045$, $df = 1$, $p = 0.832$ for females; $X^2 = 0.02368$, $df = 1$, $p = 0.878$ for males).

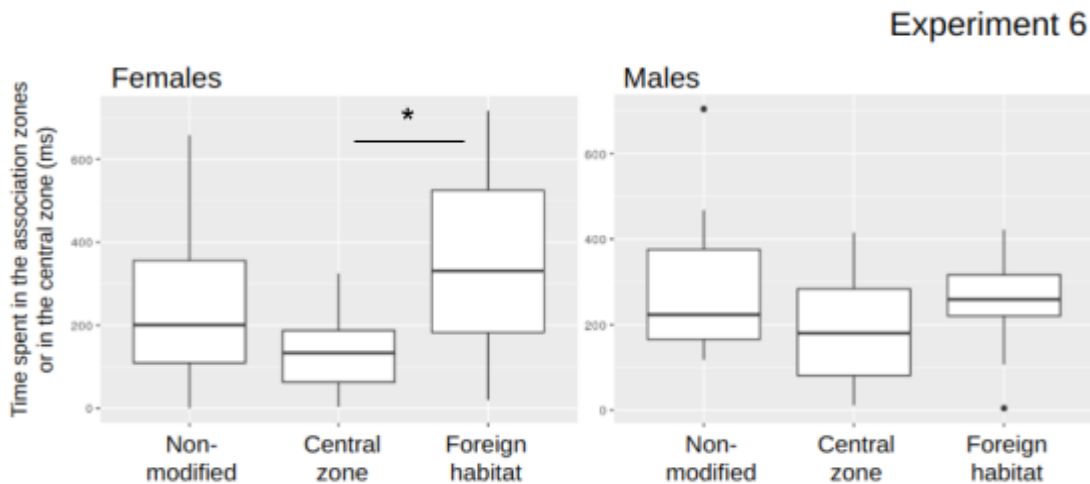


Figure 9: Time (ms) spent in the association zones of the non-modified fish or of the fish with foreign habitat statistics and in the central zones for females (left) and males (right). The video background depicts a foreign habitat (sand).

Experiment 7. Habitat preference

A processing bias hypothesis assumes that animals prefer the visual statistics of their own habitat. To estimate preference for species-typical versus foreign habitat, we analysed the data of 13 females (out of 15 tested) and 16 males (out of 20 tested). Discarded individuals either froze or stayed in the same location without further moving for the entire duration of the experiment. T-tests for the remaining males and females showed no difference between times spent on either side of the association zone (females: $t = -1.0095$, $df = 23.992$, $p = 0.3228$; males: $t = 0.79333$, $df = 29.307$, $p = 0.434$); thus we did not exclude any individuals due to side bias.

ANOVA results were significant for both males ($F(2,45) = 6.35$, $p = 0.0037$) and females ($F(2,36) = 17.3$, $p < 0.001$). Tukey's HSD Test for multiple comparisons found that the amount of time males spent in the association zone of their own habitat was significantly different from the amount of time spent in the central zone ($t = -3.54$, $p = 0.0026$). Males spent more time close to the video of their own habitat compared to the central zone.

For females, statistical differences were found between times spent close to their habitat compared to the central zone ($t = -5.85$, $p < 0.001$) and between times spent close to their habitat compared to another species' habitat ($t = -3.46$, $p = 0.0039$). Thus **females preferred their own habitat over the foreign habitat.**

On average, males entered the association zone of their own habitat 6.8 times and of the foreign habitat 6.5 times. Females entered the association zone of their own habitat 5.4 times and of the foreign habitat 4.0 times. No significant difference was found for either group ($X^2 = 0.2043$, $df = 1$, $p = 0.651$ for females; $X^2 = 0.007336$, $df = 1$, $p = 0.932$ for males).




Figure 10: Time (ms) spent in the association zones of a foreign habitat or of their own habitat and in the central zones for females (left) and males (right). The animated fish is a non-modified one.

Table 2: Summary of the significant results

	Significant results	Shown in
Experiment 1	species-specific zone > central zone	Females only
Experiment 2	species-specific zone > central zone foreign zone > central zone	Females only
Experiment 3	species-specific zone > central zone species-specific zone > non-modified	Females & males Males only
Experiment 4	species-specific zone > central zone non-modified zone > central zone	Females only Females & males
Experiment 5	species-specific zone > central zone non-modified zone > central zone	Females & males Females only
Experiment 6	foreign > central zone	Females only
Experiment 7	species-specific > central zone species-specific > foreign	Females & males Females only

Discussion

Across all experiments, females spent more time in the association zones than in the central zone, with a bias for fish modified to reflect the species-specific habitat, especially when the background video was depicting their own habitat (experiments 1 and 3). On the contrary, males spent more time in the association zones in half of the experiments only (experiments 3, 4, 5). However, when they showed a preference, males preferred animated fish modified

to reflect their own habitat over the non-modified animated fish, spending more time in the corresponding association zone than in the two other zones (experiment 3).

As for habitat preferences alone, females preferred the video of their own habitat, spending more time in the corresponding association zone than in the two other zones. Males, on the other hand, only preferred their own habitat over the central zone. For males, no difference was found between times spent alongside their own habitat versus the foreign habitat for males.

Overall, our findings suggest that both males and females tend to be sensitive to individuals modified to reflect habitat-specific statistics in their preferences, providing some support for the sensory drive theory and the processing bias hypothesis. Those biased preferences appeared more consistently for females than for males. Nonetheless, males showed an interesting result, with a preference for the fish modified to reflect the species-specific habitat over a non-modified fish. This latter result may highlight the origin of latent aesthetic preferences, which are preferences for yet-to-exist traits. Latent aesthetic preferences for complex stimuli, such as the preference for artificial crests in Australian finches (Burley & Symanski (1998), are often assumed to emerge incidentally. Our males' preferences suggest that such latent preferences might instead be determined by their visual environment and reflect a sensory bias, thus supporting a sensory drive origin of latent preferences (Endler, 1992; Endler & Basolo, 1998; Seehausen et al., 2008).

In our experiments, we kept the colour of our animated fish constant. One important future direction will be to test the interaction between colours and patterns while modulating spatial statistics of signal designs and ask how this influences female and male preferences. An interaction between colour and pattern could also bear on the interplay between detectability and camouflage. Because of their conspicuous coloration, males are more visible during the mating season, thus more closely matching natural statistics could help to make them less conspicuous to predators. This trade-off between colourful designs and camouflage has been documented in the damselfish (Phillips et al., 2017).

One way to investigate the presence of such trade-off in darters would be to test a predator of darters using the same paradigm and ask whether they attack fish that are closer to their own habitat's statistics, potentially better camouflaged, or closer to another habitat. Fish that exhibit body patterns matching the statistics of their habitat will be better camouflaged, thus less likely detected by a predator, decreasing their susceptibility to attacks.

Another future direction will be to focus on specific body parts of the fish rather than the body as a whole to assess whether some parts are more susceptible to sexual selection. For instance, Gumm and Mendelson (2011) found no correlation between body colours and fin colours of several darter species, suggesting this might also be the case for other features. This has important consequences since dorsal fins are often erected for courtship purposes or in male-male fights (see e.g. Johnston, 1994). Given that dorsal fins can be concealed, unlike the body, they may be under less predation pressure. Therefore, fin patterns may be evolving due to sexual selection alone and be less constrained by camouflage.

Conveniently, the most recent implementation of the NST algorithm (Gatys, 2017) has two control options that make it very relevant for addressing such questions: spatial and colour controls. The spatial control lets the user pre-define which part of an image to transfer the style. This is useful for modifying subparts of a body and testing which one affects visual

preferences or communication. With the colour control, the model gives the possibility to either keep the original color of the content image to only “transfer” the achromatic information or transfer both chromatic and achromatic information. The algorithm's versatility makes it convenient for additionally testing the influence of color on visual preferences.

Our suggestions for future studies show the many ways that the style transfer algorithm and similar techniques can be used to study the evolution of signal design. For instance, besides style transfer, generative adversarial networks (GANs) can be used creatively to study the evolution of animal signals. Exploiting the architecture of GANs, Talas and colleagues (2020) developed CamoGAN to study the gradual evolution of background camouflage. GANs are made of two networks that compete with each other: a discriminative one and a generative one. The discriminative network learns to make the difference between real data given as input to the model and fake data that the generative network produces. Over time, both networks become better at discriminating between real and fake data, or at generating fake data that looks more and more like real data. This makes CamoGAN a great option for simulating the evolutionary arms race between the camouflage of a synthetic prey and its predator, which is what the authors successfully did.

Conclusion

Our study demonstrates how artificial neural networks can be used to test hypotheses about animal signal evolution, by creating or manipulating stimuli without the need to target a specific feature. To make our case, we used style transfer, an easy-to-implement algorithm that the non-familiar reader will be able to adapt to their own questions. Our experiments revealed preferences for male patterns that were manipulated to be more similar to the statistics of a species-specific habitat. This preference was found in both females and males, although to different extents. Importantly, our results provide some support for the sensory drive hypothesis. As such, we showed that generative models are a powerful tool for addressing questions in visual ecology. Indeed, such models will allow us to move beyond the analysis of single features toward a more holistic perceptual representation of complex signals, paving the way to exciting new questions.

Acknowledgements

We would like to thank the students of the DART lab who helped with trials and video scoring. This work was supported by the National Science Foundation grant NSF IOS 2026334.

Supplementary material

Supplementary table: Assessing potential side bias in the six mate preference experiments

	Males	Females
Experiment 1	$t = 1.765, df = 38, p = 0.0855$	$t = -1.758, df = 31.81, p = 0.0883$
Experiment 2	$t = 1.952, df = 26.82, p = 0.0614$	$W = 94, p = 0.207$
Experiment 3	$t = 0.9881, df = 35.97, p = 0.33$	$t = 0.07073, df = 23.98, p = 0.944$
Experiment 4	$t = -0.3773, df = 26.28, p = 0.709$	$t = -1.626, df = 27.84, p = 0.115$
Experiment 5	$t = 1.071, df = 37.92, p = 0.291$	$W = 135, p = 0.76$
Experiment 6	$t = 0.2342, df = 34.42, p = 0.816$	$t = 0.834, df = 29.92, p = 0.411$

References

<https://doi.org/10.1242/jeb.151183>

Community, B. O. (2018). *Blender - a 3D modelling and rendering package*. Stichting Blender Foundation, Amsterdam. Retrieved from <http://www.blender.org>

Aspbury, A. S., & Basolo, A. L. (2002). Repeatable female preferences, mating order and mating success in the poeciliid fish, *Heterandria formosa*. *Behavioral Ecology and Sociobiology*, *51*(3), 238–244. <https://doi.org/10.1007/s00265-001-0443-1>

Brooks, R., & Endler, J. A. (2001). Female Guppies Agree to Differ: Phenotypic and Genetic Variation in Mate-Choice Behavior and the Consequences for Sexual Selection. *Evolution*, *55*(8), 1644–1655. <https://doi.org/10.1111/j.0014-3820.2001.tb00684.x>

Burley, N. T., & Symanski, Richard. (1998). “A Taste for the Beautiful”: Latent Aesthetic Mate Preferences for White Crests in Two Species of Australian Grassfinches. *The American Naturalist*, *152*(6), 792–802. <https://doi.org/10.1086/286209>

Caves, E. M., Sutton, T. T., & Johnsen, S. (2017). Visual acuity in ray-finned fishes correlates with eye size and habitat. *Journal of Experimental Biology*, *220*(9), 1586–1596. <https://doi.org/10.1242/jeb.151183>

Endler, J. A. (1992). Signals, Signal Conditions, and the Direction of Evolution. *The American Naturalist*, *139*, S125–S153. <https://doi.org/10.1086/285308>

Endler, J. A., & Basolo, A. L. (1998). Sensory ecology, receiver biases and sexual selection. *Trends in Ecology & Evolution*, *13*(10), 415–420. [https://doi.org/10.1016/s0169-5347\(98\)01471-2](https://doi.org/10.1016/s0169-5347(98)01471-2)

Fitch, H. S., & Hillis, D. M. (1984). The Anolis Dewlap: Interspecific Variability and Morphological Associations with Habitat. *Copeia*, *1984*(2), 315–323. <https://doi.org/10.2307/1445187>

Fleishman, L. J. (1992). The Influence of the Sensory System and the Environment on Motion Patterns in the Visual Displays of Anoline Lizards and Other Vertebrates. *The American Naturalist*, *139*, S36–S61. <https://doi.org/10.1086/285304>

- Fleishman, L. J., Perez-Martinez, C. A., & Leal, M. (2022). Can Sensory Drive Explain the Evolution of Visual Signal Diversity in Terrestrial Species? A Test with Anolis Lizards. *The American Naturalist*, 000–000. <https://doi.org/10.1086/720267>
- Friard, O., & Gamba, M. (2016). BORIS: A free, versatile open-source event-logging software for video/audio coding and live observations. *Methods in Ecology and Evolution*, 7(11), 1325–1330. <https://doi.org/10.1111/2041-210X.12584>
- Gatys, L. A., Ecker, A. S., & Bethge, M. (2015). A Neural Algorithm of Artistic Style. *ArXiv:1508.06576 [Cs, q-Bio]*. <http://arxiv.org/abs/1508.06576>
- Gatys, L. A., Ecker, A. S., Bethge, M., Hertzmann, A., & Shechtman, E. (2017). *Controlling Perceptual Factors in Neural Style Transfer*. 3985–3993. https://openaccess.thecvf.com/content_cvpr_2017/html/Gatys_Controlling_Perceptual_Factors_CVPR_2017_paper.html
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Networks. *ArXiv:1406.2661 [Cs, Stat]*. <http://arxiv.org/abs/1406.2661>
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144. <https://doi.org/10.1145/3422622>
- Ho, J., Jain, A., & Abbeel, P. (2020). *Denosing Diffusion Probabilistic Models* (arXiv:2006.11239). arXiv. <https://doi.org/10.48550/arXiv.2006.11239>
- Hulse, S. V., Renoult, J. P., & Mendelson, T. C. (2020). Sexual signaling pattern correlates with habitat pattern in visually ornamented fishes. *Nature Communications*, 11(1), Article 1. <https://doi.org/10.1038/s41467-020-16389-0>
- Johnston, C. E. (1994). Spawning Behavior of the Goldstripe Darter (*Etheostoma parvipinne* Gilbert and Swain) (Percidae). *Copeia*, 1994(3), 823–825. <https://doi.org/10.2307/1447204>
- Kingma, D. P., & Welling, M. (2019). An Introduction to Variational Autoencoders. *Foundations and Trends® in Machine Learning*, 12(4), 307–392.

<https://doi.org/10.1561/22000000056>

- Lehtonen, T. K., & Lindström, K. (2008). Repeatability of mating preferences in the sand goby. *Animal Behaviour*, 75(1), 55–61. <https://doi.org/10.1016/j.anbehav.2007.04.011>
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., & Chen, M. (2022). *GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models* (arXiv:2112.10741). arXiv. <https://doi.org/10.48550/arXiv.2112.10741>
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., ... Chintala, S. (2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library* (arXiv:1912.01703). arXiv. <https://doi.org/10.48550/arXiv.1912.01703>
- Pérez-Rodríguez, L., Jovani, R., & Stevens, M. (2017). Shape matters: Animal colour patterns as signals of individual quality. *Proceedings of the Royal Society B: Biological Sciences*, 284(1849), 20162446. <https://doi.org/10.1098/rspb.2016.2446>
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., & Sutskever, I. (2021). *Zero-Shot Text-to-Image Generation* (arXiv:2102.12092). arXiv. <http://arxiv.org/abs/2102.12092>
- Renoult, J. P., & Mendelson, T. C. (2019). Processing bias: Extending sensory drive to include efficacy and efficiency in information processing. *Proceedings of the Royal Society B: Biological Sciences*, 286(1900), 20190165. <https://doi.org/10.1098/rspb.2019.0165>
- Seehausen, O., Terai, Y., Magalhaes, I. S., Carleton, K. L., Mrosso, H. D. J., Miyagi, R., van der Sluijs, I., Schneider, M. V., Maan, M. E., Tachida, H., Imai, H., & Okada, N. (2008). Speciation through sensory drive in cichlid fish. *Nature*, 455(7213), Article 7213. <https://doi.org/10.1038/nature07285>
- Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., & Ganguli, S. (2015). *Deep Unsupervised Learning using Nonequilibrium Thermodynamics* (arXiv:1503.03585).

arXiv. <https://doi.org/10.48550/arXiv.1503.03585>

Song, Y., & Ermon, S. (2020). *Generative Modeling by Estimating Gradients of the Data Distribution* (arXiv:1907.05600). arXiv. <https://doi.org/10.48550/arXiv.1907.05600>

Talas, L., Fennell, J. G., Kjærsmo, K., Cuthill, I. C., Scott-Samuel, N. E., & Baddeley, R. J. (2020). CamoGAN: Evolving optimum camouflage with Generative Adversarial Networks. *Methods in Ecology and Evolution*, 11(2), 240–247. <https://doi.org/10.1111/2041-210X.13334>

Wilson, A. J., Dean, M., & Higham, J. P. (2013). A game theoretic approach to multimodal communication. *Behavioral Ecology and Sociobiology*, 67(9), 1399–1415. <https://doi.org/10.1007/s00265-013-1589-3>

Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2020). Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *ArXiv:1703.10593 [Cs]*. <http://arxiv.org/abs/1703.10593>