

Emergent intensity invariance vs. signal-to-noise ratio at three consecutive processing stages along the grasshopper song recognition pathway

Jona Hartling, Jan Benda

1 Introduction

Intensity invariance is a fundamental property of sensory systems across modalities and species, from fruit flies (Ozeri-Engelhard et al. 2018; Ketkar et al. 2023) over crickets (Benda and R Matthias Hennrich 2018) and grasshoppers (Clemens, Weschke, et al. 2010) to rodents (Bolding and Franks 2018) and primates (Barbour 2011). It allows for the robust recognition of behaviorally relevant stimuli despite variations in stimulus intensity. However, the computational mechanisms underlying intensity invariance are often difficult to disentangle. Here, we use a physiologically inspired functional model of the grasshopper song recognition pathway to investigate the emergence of intensity invariance throughout the auditory processing stream.

The auditory system of grasshoppers (*Acrididae*) has been studied extensively over the years. Grasshoppers rely on their sense of hearing for intraspecific communication — including mate attraction (D. v. Helversen 1972) and evaluation (Stange and Ronacher 2012), sender localization (D. v. Helversen and Rheinlaender 1988), courtship display (Elsner 1968), and rival deterrence (Greenfield and Minckley 1993) — and have evolved a variety of acoustic signals for different behavioral contexts (Otte 1970). The most conspicuous acoustic signals of grasshoppers are their species-specific calling songs, which broadcast the presence of the singing individual to potential mates within range. These songs are usually more characteristic of a species than morphological traits (Tishechkin and V Yu Vedenina 2016; Tarasova et al. 2021), which can vary greatly within species (Rowell 1972; Köhler et al. 2017). The reliance on songs to mediate reproduction represents a strong evolutionary driving force that resulted in a massive species diversification (Varvara Yu Vedenina and Mugue 2011; Sevastianov et al. 2023), with over 6800 recognized species in the *Acrididae* family (Cigliano et al. 2024).

Grasshopper songs are amplitude-modulated broad-band acoustic signals. They consist of a series of noisy syllables and relatively quiet pauses, which form a characteristic repetitive pattern (Otto von Helversen and Elsner 1977; Stumpner and Otto von Helversen 1994). Song recognition depends on certain structural parameters of this pattern — such as the ratio of syllable duration to pause duration (D. v. Helversen 1972), the slope of pulse onsets (D. v. Helversen 1993), and the accentuation of syllable onsets relative to the preceding pause (Balakrishnan et al. 2001; D. v. Helversen, Balakrishnan, and Helversen 2004) — which are sufficiently conveyed by the amplitude modulation of the song alone (D. v. Helversen and Otto von Helversen 1997).

Grasshopper songs, like all acoustic signals, are subject to sound attenuation, which depends on the distance from the sound source, the frequency content of the signal, and the vegetation of the habitat (Michelsen 1978). Sound attenuation has two major consequences for song recognition. First, the amplitude dynamics of the song pattern degrade with increasing distance to the sender, which limits the effective communication range of grasshoppers to 1-2 m in their typical grassland habitats (Lang 2000). Second, the intensity of a song at the receiver’s position varies with the position of the sender, which should ideally not affect song recognition. The auditory system thus needs to achieve a certain degree of intensity invariance — a time scale-selective sensitivity to faster amplitude dynamics and simultaneous insensitivity to more sustained amplitude dynamics. Intensity invariance is commonly associated with neural adaptation (Benda and R Matthias Hennig 2008; Barbour 2011; Ozeri-Engelhard et al. 2018; more general: Benda 2021). Different neuron types in the grasshopper auditory system exhibit spike-frequency adaptation in response to sustained stimulation (Römer 1976; Gollisch and Herz 2004; Hildebrandt et al. 2009; Clemens, Weschke, et al. 2010; Fisch et al. 2012). Accordingly, intensity invariance is not the result of a single processing step but rather a gradual process, in which different neuronal populations contribute to varying degrees (Clemens, Weschke, et al. 2010) and by different mechanisms (Hildebrandt et al. 2009).

In the current study, we leverage functional modelling to trace the emergence of intensity invariance through individual processing steps of the grasshopper song recognition pathway. The model pathway we propose here is based on a previous functional model framework for song recognition in both crickets (Clemens and R Matthias Hennig 2013; R Matthias Hennig et al. 2014) and grasshoppers (Clemens and Bernhard Ronacher 2013; review on both: Bernhard Ronacher, R Matthias Hennig, and Clemens 2015). The existing framework relies on pulse trains as input signals, which were designed to capture the essential structural properties of natural song envelopes (Clemens and Bernhard Ronacher 2013). It includes feature extraction by a bank of linear-nonlinear feature detectors, evidence accumulation by temporal averaging of each feature, and categorical decision making by a weighted linear combination of feature values. We adopted the general structure of the existing framework and extended it by a physiologically plausible preprocessing stage — including spectral filtering, envelope extraction, logarithmic compression, and intensity adaptation — which allows the model to operate on unmodified recordings of natural grasshopper songs. The resulting model pathway thus covers the entire auditory processing stream from the initial reception of airborne sound waves to the generation of a high-dimensional feature representation that allows for the categorical recognition of conspecific songs. It incorporates anatomical, physiological, and ethological evidence from several decades of research on the grasshopper auditory system. In the following, we provide a side-by-side account of the known physiological processing steps along the song recognition pathway and their functional approximations in the model pathway. We then elaborate on the computational mechanisms that contribute to the emergence of intensity-invariant song representations, the interaction between these mechanisms, the overall capacity for intensity invariance in the system, and the ethological implications of our findings.

2 Methods

All modeling, data analysis, and data visualization was performed in Python 3.12.3 except for the pathway overview (Fig. 1), which was assembled in Inkscape 1.2. The code base for the model pathway is available as the *thunderhopper* package, version 1.0, on PyPi. Any audio data was inspected and edited with the help of the *audian* package, version 2.4, on PyPi.

2.1 Functional model of the grasshopper song recognition pathway

The anatomical organisation of the grasshopper song recognition pathway can be outlined as a feed-forward network of three consecutive neuronal populations (Fig. 1a-c): Peripheral auditory receptor neurons, whose axons enter the ventral nerve cord (VNC) at the level of the metathoracic ganglion; local interneurons that remain exclusively within the thoracic region of the VNC; and ascending neurons projecting from the thoracic region towards the supraesophageal ganglion (SEG), or central brain (Rehbein et al. 1974; Rehbein 1976; Eichendorf and Kalmring 1980). The input to the network originates at the tympanal membrane, which acts as acoustic receiver and is coupled to the dendritic endings of the receptor neurons (Gray 1960). The outputs from the network converge in the SEG, which presumably harbors the neuronal substrate for conspecific song recognition and response initiation (Bernhard Ronacher, D. v. Helversen, and Helversen 1986; Bauer and Helversen 1987; Bhavsar et al. 2017).

Around 15 to 20 ascending neurons have been identified in the grasshopper auditory system (Stumpner and Bernhard Ronacher 1991), whose functional characteristics are conserved even between species that are not closely related (Neuhofer et al. 2008). The population of ascending neurons possesses a diverse range of response properties that contrasts with the rather homogeneous responses of receptor neurons and local interneurons (Clemens, Kutzki, et al. 2011), which suggests a transition from a uniform population-wide processing stream into several parallel branches. Accordingly, the model pathway is divided into two distinct stages (Fig. 1d): The preprocessing stage incorporates the processing steps at the levels of the tympanal membrane, the receptor neurons, and the local interneurons; and operates on one-dimensional signal representations (Fig. 2). The feature extraction stage corresponds to the processing within the ascending neurons and further downstream towards the SEG; and operates on high-dimensional signal representations (Fig. 3). The details of each physiological processing step and its functional approximation are described in the following sections.

Around 15 to 20 ascending neurons have been identified in the grasshopper auditory system (Stumpner and Bernhard Ronacher 1991), whose functional characteristics are conserved even between species that are not closely related (Neuhofer et al. 2008). The population of ascending neurons possesses a diverse range of response properties that contrasts with the rather homogeneous responses of receptor neurons and local interneurons (Clemens, Kutzki, et al. 2011), which suggests a transition from a uniform population-wide processing stream into several parallel branches. Accordingly, the model pathway is divided into two distinct stages (Fig. 1d): The preprocessing stage incorporates the processing steps at the levels of the tympanal membrane, the receptor neurons, and the local interneurons; and operates on one-dimensional signal representations (Fig. 2). The feature extraction stage corresponds to the processing within the ascending neurons and further downstream towards the SEG; and operates on high-dimensional signal representations (Fig. 3). The details of each physiological processing step and its func-

tional approximation are described in the following sections.

2.1.1 Population-driven signal preprocessing

Grasshoppers receive airborne sound waves by a tympanal organ at each side of the body. The tympanal membrane acts as a mechanical resonance filter for sound-induced vibrations (Windmill et al. 2000; Malkin et al. 2014). Vibrations that fall within specific frequency bands are focused on different membrane areas, while others are attenuated. This processing step can be approximated by an initial bandpass filter (Fig. 2a) applied to the acoustic input signal $x_{\text{raw}}(t)$:

$$x_{\text{filt}}(t) = x_{\text{raw}}(t) * h_{\text{BP}}(t), \quad f_{\text{cut}} = 5 \text{ kHz}, 30 \text{ kHz} \quad (1)$$

The receptor neurons transduce the vibrations of the tympanal membrane into sequences of action potentials. They thereby encode the amplitude modulation, or envelope, of the signal (Machens, Prinz, et al. 2001), which likely involves a rectifying nonlinearity (Machens, Stemmler, et al. 2001). The extraction of the signal envelope (Fig. 2b) can be modelled as full-wave rectification followed by lowpass filtering of the tympanal signal $x_{\text{filt}}(t)$:

$$x_{\text{env}}(t) = |x_{\text{filt}}(t)| * h_{\text{LP}}(t), \quad f_{\text{cut}} = 250 \text{ Hz} \quad (2)$$

Furthermore, the receptors exhibit a sigmoidal response curve over logarithmically compressed stimulus intensities (Suga 1960; Gollisch, Schütze, et al. 2002). In the model pathway, logarithmic compression (Fig. 2c) is achieved by conversion to decibel scale

$$x_{\text{log}}(t) = 20 \cdot \log_{10} \frac{x_{\text{env}}(t)}{x_{\text{ref}}}, \quad x_{\text{ref}} = 1 \quad (3)$$

relative to the common reference intensity x_{ref} . Both the receptor neurons (Römer 1976; Gollisch and Herz 2004; Fisch et al. 2012) and, on a larger scale, the subsequent local interneurons (Hildebrandt et al. 2009; Clemens, Weschke, et al. 2010) adapt their firing rates in response to sustained stimulus intensities, which allows for the robust encoding of faster amplitude modulations against a slowly changing overall baseline intensity. Functionally, the adaptation mechanism resembles a highpass filter (Fig. 2d) over the logarithmically compressed envelope $x_{\text{log}}(t)$:

$$x_{\text{adapt}}(t) = x_{\text{log}}(t) * h_{\text{HP}}(t), \quad f_{\text{cut}} = 10 \text{ Hz} \quad (4)$$

This processing step concludes the preprocessing stage of the model pathway. The resulting intensity-adapted envelope $x_{\text{adapt}}(t)$ is then passed on from the local interneurons to the ascending neurons, where it serves as the basis for the following feature extraction stage.




Fig. 1: Schematic organisation of the grasshopper song recognition pathway and structure of the functional model pathway. **a:** Simplified course of the pathway in the grasshopper, from the tympanal membrane over receptor neurons, local interneurons, and ascending neurons further towards the supraesophageal ganglion. **b:** Schematic of synaptic connections between the three neuronal populations within the metathoracic ganglion. **c:** Network representation of neuronal connectivity. **d:** Flow diagram of consecutive signal representations (boxes) and transformations (arrows) along the model pathway. All representations are time-varying. 1st half: Preprocessing stage (one-dimensional representations). 2nd half: Feature extraction stage (high-dimensional representations).




Fig. 2: Song representations during the preprocessing stage. Example song of *O. rufipes*. **a:** Bandpass filtered tympanal signal $x_{\text{filt}}(t)$. **b:** Signal envelope $x_{\text{env}}(t)$. **c:** Logarithmically compressed envelope $x_{\text{log}}(t)$. **d:** Intensity-adapted envelope $x_{\text{adapt}}(t)$.

2.1.2 Feature extraction by individual neurons

The population of ascending neurons extracts and encodes a number of different features of the preprocessed signal, and hence represents the signal in a higher-dimensional space than the preceding receptor neurons and local interneurons (Clemens, Kutzki, et al. 2011). Each ascending neuron is assumed to scan the signal for a specific template pattern, which can be thought of as a kernel of a particular structure and on a particular time scale. This process, known as template matching, can be modelled as a convolution of the intensity-adapted envelope $x_{\text{adapt}}(t)$ with a kernel $k_i(t)$ specific to the i -th ascending neuron:

$$c_i(t) = x_{\text{adapt}}(t) * k_i(t) = \int_{-\infty}^{+\infty} x_{\text{adapt}}(\tau) \cdot k_i(t - \tau) d\tau \quad (5)$$

We use Gabor kernels as basis functions for creating different template patterns. Gabor functions presumably capture the essential structural properties of the filter functions found in various auditory neurons (Rokem et al. 2006; Clemens, Kutzki, et al. 2011; Clemens, Wohlgemuth, and Ronac). An arbitrary one-dimensional, real Gabor kernel is generated by multiplication of a Gaussian envelope with standard deviation or kernel width σ_i and a sinusoidal carrier with frequency ω_i and phase ϕ_i :

$$k_i(t, \sigma_i, \omega_i, \phi_i) = e^{-\frac{t^2}{2\sigma_i^2}} \cdot \sin(\omega_i t + \phi_i), \quad \omega_i = 2\pi f_{\text{sin}_i} \quad (6)$$

Different combinations of σ_i and ω_i result in Gabor kernels with different lobe number n_i , which is the number of half-periods of the carrier that fit under the Gaussian envelope within reasonable limits of attenuation. The time window under the Gaussian envelope that contains

the relevant lobes of the kernel can be defined as Gaussian full duration at height h_{rel} relative to the maximum of the Gaussian:

$$\text{FDRM}(\sigma_i, h_{\text{rel}}) = 2 \cdot \sqrt{-2 \cdot \ln h_{\text{rel}}} \cdot \sigma_i, \quad h_{\text{rel}} \in (0, 1] \quad (7)$$

With this, an appropriate carrier frequency ω_i for obtaining a Gabor kernel with width σ_i and desired lobe number n_i can be approximated as

$$\omega_i(n_i, \sigma_i, h_{\text{rel}}) = \frac{0.5 \cdot n_i + \beta_0}{\text{FDRM}(\sigma_i, h_{\text{rel}})}, \quad n_i \geq 2 \forall n_i \in \mathbb{Z} \quad (8)$$

The relationship between ω_i and n_i is approximately linear except for small n_i . The offset term $\beta_0 \approx 0.26$ was added to balance the amplitudes of the n_i desired lobes of the kernel — which should be maximized — against the amplitudes of the next-outer lobes, which should not exceed the threshold value determined by h_{rel} . Note that simple Gaussian kernels with $n_i = 1$ can be obtained by setting the carrier frequency to $\omega_i = 0$ and are hence not covered by Eq. 8.

Carrier phase ϕ_i determines the position of the kernel lobes relative to the kernel center. We restrict the Gabor kernels to be either even or odd functions by setting ϕ_i to one of only four specific phase values (Tab. 1). Even Gabor kernels are mirror-symmetric with uneven n_i , whereas odd Gabor kernels are point-symmetric with even n_i . Both even and odd kernels can have either positive or negative sign, which refers to the sign of the kernel’s central lobe (even kernels) or the left of the two central lobes (odd kernels). These four major groups of Gabor kernels allow for the extraction of different types of signal features, such as the presence of peaks (even, +), troughs (even, -), onsets (odd, +), and offsets (odd, -) at various time scales.

Tab. 1: Values of phase ϕ that are specific for the four major groups of Gabor kernels.

sign	even kernels	odd kernels
+	$+\pi/2$	π
-	$-\pi/2$	0

Following the convolutional template matching (Fig. 3a), each kernel-specific response $c_i(t)$ is passed through a shifted Heaviside step-function $H(c_i - \Theta_i)$ with threshold value Θ_i to obtain a binary response (Fig. 3b):

$$b_i(t, \Theta_i) = \begin{cases} 1, & c_i(t) > \Theta_i \\ 0, & c_i(t) \leq \Theta_i \end{cases} \quad (9)$$

The thresholding of $c_i(t)$ into $b_i(t)$ can be thought of as a categorization into ”relevant” and ”irrelevant” response values. Similar thresholding nonlinearities have been a crucial processing step in previous models that deal with the extraction of behaviorally relevant song features in insects (Clemens and R Matthias Hennig 2013; Clemens and Bernhard Ronacher 2013; R Matthias Hennig Bernhard Ronacher, R Matthias Hennig, and Clemens 2015).

In the grasshopper, the responses of the ascending neurons are assumed to be integrated somewhere in the SEG (Bernhard Ronacher, D. v. Helversen, and Helversen 1986; Bauer and Helversen 1987;

Bhavsar et al. 2017). In the model pathway, temporal integration is implemented as temporal averaging of the binary responses $b_i(t)$ by a lowpass filter with extremely low cutoff frequency:

$$f_i(t) = b_i(t) * h_{LP}(t), \quad f_{cut} = 1 \text{ Hz} \quad (10)$$

This processing step results in a set of slowly changing kernel-specific features $f_i(t)$, which is the final representation along the model pathway (Fig. 3c). In the resulting high-dimensional feature space, different species-specific song patterns can be distinguished by their distinct combination of feature values, e.g. using Euclidian geometry or a simple linear classifier.




Fig. 3: Song representations during the feature extraction stage. Example song of *O. rufipes*. Different color shades indicate different types of Gabor kernels with specific lobe number n_i and either + or - sign, sorted (dark to light) first by increasing n_i and then by sign ($1 \leq n_i \leq 4$; first +, then - for each n_i ; two kernel widths σ_i of 4 ms and 32 ms per type; 8 types, 16 kernels in total). **a:** Kernel-specific filter responses $c_i(t)$. **b:** Binary responses $b_i(t)$. **c:** Finalized features $f_i(t)$.

2.2 Simulation-based analysis of the model pathway

2.2.1 Data sourcing

All simulations were based on a dataset that was assembled from five different sources, each of which is an established reference for the identification of European grasshopper species. The dataset was limited to six species from the species-rich *Gomphocerinae* sub-family that are known to be common throughout Central and Southern Europe. All recordings were converted to standard *.wav* format with a sampling rate of 44.1 kHz and an amplitude scale in arbitrary units. Individual songs were then cut from each recording. The dataset includes a total of 31 recordings across species, which amounts to a total of 153 isolated songs. However, the number of available species-specific songs varies greatly across species, with a maximum of 48 songs for *C. biguttulus* and a minimum of 6 songs for *C. mollis* (Tab. 2).

- "Heuschrecken beobachten, bestimmen" by Heiko Bellmann
1st edition, 1993, Naturbuch, Augsburg
- "Gesänge der heimischen Heuschrecken. Akustisch-optische Bestimmungshilfe."
by Karl-Heinz Garberding, Deutscher Jugendbund für Naturbeobachtung
1st edition, 2001, DJN, Göttingen
- "Heuschrecken – Die Stimmen von 61 heimischen Arten" by Heiko Bellmann
1st edition, 2004, AMPLE, Germering
- "Fauna d'Italia XLVIII – Orthoptera" by Bruno Massa, Paolo Fontana, Filippo M. Buzzetti,
Roy M.J.C. Kleukers, Baudewijn Odé
1st edition, 2012, edagricola, Milano
- "Singing Orthoptera of Slovenia" by Stanislav Gomboc, Blaz Segula
1st edition, 2014, EGEA, Ljubljana

Tab. 2: Overview of the six grasshopper species from the *Gomphocerinae* sub-family, the number of sources per species, the number of available recordings across sources, and the number of isolated songs across recordings.

Species	Sources	Recordings	Songs
<i>Chorthippus biguttulus</i>	5	6	48
<i>Chorthippus mollis</i>	3	3	6
<i>Chrysochraon dispar</i>	4	5	45
<i>Gomphocerippus rufus</i>	4	8	16
<i>Omocestus rufipes</i>	4	5	14
<i>Pseudochorthippus parallelus</i>	4	4	24

2.2.2 Generation of synthetic input signals

Different processing steps along the model pathway were tested for intensity invariance by generating synthetic input signals $x(t)$ of varying intensity, transforming them through the respective processing steps, and comparing the resulting signal representations. Inputs were generated for two distinct cases. In the idealized, noiseless case, $x(t)$ consists of a song component $s(t)$ with $\sigma_s = 1$ and a multiplicative scale α :

$$x(t) = \alpha \cdot s(t), \quad \alpha \geq 0 \quad (11)$$

In the noiseless case, $x(t)$ is hence only a scaled version of $s(t)$ with $\sigma_x = \alpha$. In the more realistic, noisy case, $x(t)$ consists of the same song component $s(t)$ scaled by α and an additive noise component $\eta(t)$ with $\sigma_\eta = 1$:

$$x(t) = \alpha \cdot s(t) + \eta(t), \quad \alpha \geq 0 \quad (12)$$

Accordingly, the signal-to-noise ratio (SNR) of input $x(t)$ in the noisy case equals the squared α value:

$$\text{SNR}_x(\alpha) = \frac{(\alpha \cdot \sigma_s)^2}{\sigma_\eta^2} = \alpha^2, \quad \sigma_s = \sigma_\eta = 1 \quad (13)$$

For most analyses, it would be sufficient if input $x(t)$ corresponds to the signal representation immediately before the first of the tested transformations. For instance, when testing the effects of logarithmic compression (Eq. 3), $x(t)$ would correspond to the signal envelope $x_{\text{env}}(t)$. However, in this particular case, $x_{\text{env}}(t)$ results from a nonlinear transformation (Eq. 2), which cannot be synthesized as an additive mixture of $s(t)$ and $\eta(t)$. For this reason, any input $x(t)$ across all analyses corresponds not to the representation immediately before the tested transformations but its predecessor representation instead. Therefore, when testing logarithmic compression, $x(t)$ corresponds to the tympanal signal $x_{\text{flt}}(t)$ instead of $x_{\text{env}}(t)$.

The raw $s(t)$ was drawn from the dataset of isolated species-specific song recordings, whereas the raw $\eta(t)$ consists of a segment of normally distributed white noise. Both $s(t)$ and $\eta(t)$ were normalized to unit standard deviation. These can be used without further processing for all analyses where input $x(t)$ corresponds to $x_{\text{raw}}(t)$. For analyses where $x(t)$ corresponds to a later representation, $s(t)$ and $\eta(t)$ were first processed along the model pathway up to the required representation, again normalized to unit standard deviation, and then used to generate $x(t)$ according to either Eq. 11 in the noiseless case or Eq. 12 in the noisy case.

2.2.3 Quantifying signal intensity across representations

All intensity measures were calculated over a manually labeled segment within each song. Segments always excluded the first and last few syllables to allow slowly changing representations such as $f_i(t)$ to stabilize. The duration of each segment and the number of contained syllables depends on the duration of the species-specific song. Care was taken to ensure that the segment contained a sufficient number of syllables to obtain a reliable estimate of the intensity measures.

The standard deviation σ was used as a measure of intensity for all representations resulting from the transformation of input $x(t)$ up to and including the kernel responses $c_i(t)$, for which individual σ_{c_i} were used as kernel-specific intensity measures. The binary responses $b_i(t)$ were deemed to similar to the features $f_i(t)$ to warrant their own intensity measure and were hence omitted from all related analyses. For $f_i(t)$, σ is not an appropriate intensity measure because each $f_i(t)$ is ideally constant with $\sigma = 0$ for the duration of a song. Therefore, the average value μ_{f_i} of each $f_i(t)$ was used as a kernel-specific intensity measure instead.

It is arguably not ideal to quantify the intensity of $c_i(t)$ and $f_i(t)$ separately for each kernel. Overall, these representations are not separate signals bundled together but rather a set that acts as a unit with a single intensity measure. However, there is no straightforward way to quantify the intensity of $c_i(t)$ or $f_i(t)$ as a whole that would not entail a certain ambiguity, e.g by averaging across kernels. In this sense, we opted for the kernel-specific approach because it allows to asses differences in the dependency on α between individual members of either $c_i(t)$ and $f_i(t)$.

The absolute intensity measures allow to compare the intensity of a representation across different α values. Additionally, ratios were calculated between the intensity measures for $\alpha > 0$ and the respective pure-noise reference measure for $\alpha = 0$ to better compare the intensities of different representations. This is only possible in the noisy case, where input $x(t) = \eta(t)$ for $\alpha = 0$, whereas $x(t) = 0$ for $\alpha = 0$ in the noiseless case. At the level of input $x(t)$, the ratio of

intensity measures depends on the square root of α :

$$\frac{\sigma_x}{\sigma_\eta} = \sqrt{\frac{\sigma_x^2}{\sigma_\eta^2}} = \sqrt{\frac{(\alpha \cdot \sigma_s)^2 + \sigma_\eta^2}{\sigma_\eta^2}} = \sqrt{\alpha^2 + 1}, \quad \sigma_s = \sigma_\eta = 1 \quad (14)$$

This holds only if $s(t) \perp \eta(t)$, so that $\sigma_x^2 = \sigma_s^2 + \sigma_\eta^2$, which is a reasonable assumption for the raw $s(t)$ and $\eta(t)$. However, the dependency of the ratio on α is not necessarily the same for representations that are transformed from $x(t)$ by nonlinear operations, since these change the relationship of $s(t)$ and $\eta(t)$ in an unpredictable fashion (see appendix Fig. 12). Furthermore, the ratio is not a proper SNR of the representation because it does not relate $s(t)$ to $\eta(t)$ within the representation but rather the entire representation to $\eta(t)$ alone. However, it still provides a useful measure of the relative intensity of a representation with and without $s(t)$, which is the closest we can get to the SNR of the representation. As such, the ratio of intensity measures is referred to as SNR in the following.

2.3 Field data-based analysis of the model pathway

Field recordings were taken on a meadow in the vicinity of the University of Tübingen, Germany, during the day in August 2024. All recordings were taken using a custom hand-held microphone array that was assembled from eight omnidirectional AV-TEFE TCM141 condenser microphones. The microphones were arranged in a linear configuration with a spacing of 30 cm between adjacent microphones and oriented in the same direction along the axis of the array. All microphones were connected to a custom 8-channel amplification and digitization system based on a Teensy 4.1 microcontroller with real-time clock and microSD card storage. Recordings were written to the microSD card in *.wav* format with a sampling rate of 96 kHz and an amplitude scale in arbitrary units. The microphone array was held at a height of approximately 30 cm above the ground, which was slightly above the height of most surrounding vegetation and at the same height as the singing grasshopper. The array was moved as close to the grasshopper as possible without interrupting its song production, which amounts to an approximate offset distance of 10 cm between the animal and the leading microphone. Care was taken to maintain a stable position and height of the microphone array during recording. The resulting recordings were then processed through the model pathway and analyzed according to the procedure described in Section 2.2.3.

2.4 Determining kernel-specific threshold values

Different kernels $k_i(t)$ result in specific kernel responses $c_i(t)$, Eq. 5, which are then transformed further into binary responses $b_i(t)$, Eq. 9, by thresholding nonlinearity $H(c_i - \Theta_i)$. The threshold value Θ_i is specific to each $k_i(t)$. Across all analyses, Θ_i has been specified as a multiple of the pure-noise reference standard deviation σ_{c_i} for input $x(t) = \eta(t)$. This ensures that Θ_i as well as the resulting $b_i(t)$ and $f_i(t)$ are comparable across different $k_i(t)$ because each pure-noise $c_i(t)$ approximately follows a normal distribution around zero (see appendix Figs. 18-21).

3 Results

3.1 Mechanisms driving the emergence of intensity invariance

It is not necessary to test each processing step along the model pathway for intensity invariance. Instead, we can focus on those steps that involve nonlinear transformations, since these are the only steps that can potentially change the dependency on scale α between the input and output representations. Overall, there are three nonlinear transformations along the model pathway: Full-wave rectification during envelope extraction, logarithmic compression, and the thresholding nonlinearity during feature extraction. In the following, we analyze the effects of each of these transformations on the intensity and SNR of the resulting representations as well as their potential contribution to intensity invariance.

3.1.1 Full-wave rectification & lowpass filtering

The first nonlinear transformation along the model pathway is the full-wave rectification of the tympanal signal $x_{\text{filt}}(t)$ during the extraction of the signal envelope (Eq. 2). Rectification transforms the distribution of $x_{\text{filt}}(t)$ from an approximately zero-centered distribution with both positive and negative values into a strictly non-negative distribution. Signal envelope $x_{\text{env}}(t)$ is then obtained by lowpass filtering the rectified $x_{\text{filt}}(t)$. The effects of this transformation pair on SNR and potential intensity invariance were analyzed by rescaling and processing the input signal $x_{\text{raw}}(t)$ and comparing standard deviations between the resulting $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$, once for the noiseless case (Fig. 4a) and once for the noisy case (Fig. 4b). In addition, the cutoff frequency f_{cut} of the lowpass filter was varied to investigate the influence of different filter bandwidths. In the noiseless case, the standard deviations of $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$ are each reduced compared to the input $x_{\text{raw}}(t)$ by a multiplicative factor. These factors are constant across all α , which results in a downward shift of the respective curve on a double-logarithmic scale, away from the diagonal (Fig. 4c). For $x_{\text{filt}}(t)$, the reduction is a consequence of the bandpass filtering (Eq. 1) of $x_{\text{raw}}(t)$. For $x_{\text{env}}(t)$, the standard deviation is further reduced compared to $x_{\text{filt}}(t)$. Rectification contributes much less to this reduction than lowpass filtering. The degree of reduction by lowpass filtering depends on the cutoff frequency f_{cut} , with lower f_{cut} (narrow bandwidth) resulting in a stronger reduction. In the noisy case, the standard deviations of $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$ can be related to the respective pure-noise reference standard deviation (Fig. 4d). This causes each curve to start with a constant regime of SNR values near 1 for smaller α , which reflects the dominance of the noise component $\eta(t)$ over the song component $s(t)$ in the input $x_{\text{raw}}(t)$. For larger α , all curves transition into a regime of linearly increasing SNR on a double-logarithmic scale. For $x_{\text{filt}}(t)$, the linear part of the curve deviates only slightly from the diagonal. For $x_{\text{env}}(t)$, however, the transition occurs at lower α compared to $x_{\text{filt}}(t)$, and the linear part of the curve is shifted leftward away from the diagonal, which means that higher SNR values are achieved for the same α . This effect is more pronounced for lower f_{cut} of the lowpass filter and is presumably caused by the attenuation of high-frequency components in the signal, which are more prominent in the noise component $\eta(t)$ than in the song component $s(t)$. The effect also appears relatively consistent across different species, although small variations exist (Fig. 4e and appendix Fig. 13). In summary, the standard deviation of $x_{\text{env}}(t)$ has never been observed to saturate for larger α but rather continues to increase proportionally to α for all tested f_{cut} , in both the noiseless and the noisy

case and across different species. Consequently, the combination of rectification and lowpass filtering does not contribute to intensity invariance. However, this transformation pair does improve the SNR of $x_{\text{env}}(t)$ relative to $x_{\text{filt}}(t)$ and thus provides subsequent processing stages with a more robust input representation and higher input SNR.




Fig. 4: Rectification and lowpass filtering improves SNR but does not contribute to intensity invariance. Input $x_{\text{raw}}(t)$ consists of $s(t)$ scaled by α with optional $\eta(t)$ and is successively transformed into tympanal signal $x_{\text{filt}}(t)$ and envelope $x_{\text{env}}(t)$. **Top:** Examples of $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$ for different α . **a:** Noiseless case. **b:** Noisy case. **Bottom:** Intensity measures over α . Different line styles indicate different cutoff frequencies f_{cut} of the lowpass filter extracting $x_{\text{env}}(t)$. **c:** Noiseless case: Standard deviation σ_x of $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$, respectively. **d:** Noisy case: Ratio of σ_x to the respective pure-noise reference σ_η for $\alpha = 0$. **e:** Ratio of σ_x to σ_η of $x_{\text{env}}(t)$ as in **d** for different species (averaged over songs and recordings, appendix Fig. 13).

3.1.2 Logarithmic compression & spike-frequency adaptation

The second nonlinear transformation along the model pathway is the logarithmic compression of the signal envelope $x_{\text{env}}(t)$ into $x_{\text{log}}(t)$, Eq. 3, which is then followed by the highpass filtering of $x_{\text{log}}(t)$, Eq. 4, to obtain the intensity-adapted envelope $x_{\text{adapt}}(t)$. The interplay of this transformation pair was analyzed by rescaling and processing the input signal $x_{\text{filt}}(t)$ and comparing standard deviations between the resulting $x_{\text{env}}(t)$, $x_{\text{log}}(t)$, and $x_{\text{adapt}}(t)$. It is necessary to use $x_{\text{filt}}(t)$ as input for this analysis instead of $x_{\text{env}}(t)$, because $x_{\text{env}}(t)$ results from a nonlinear transformation and hence cannot be synthesized as an additive mixture of song component $s(t)$ and noise component $\eta(t)$. However, it is much easier to conceive a mathematical description of

the effects of logarithmic compression and adaptation if $x_{\text{env}}(t)$ itself is assumed to be composed of $s(t)$ and $\eta(t)$. In the noiseless case (Fig. 5a), $x_{\text{env}}(t)$ takes the form of

$$x_{\text{env}}(t) = \alpha \cdot s(t), \quad x_{\text{env}}(t) > 0 \quad \forall t \in \mathbb{R} \quad (15)$$

The standard deviation of $x_{\text{env}}(t)$ increases linearly with α on a double-logarithmic scale and is slightly reduced (Fig. 5c) compared to the input $x_{\text{filt}}(t)$, which is consistent with the results of the previous analysis (Fig. 4c). By conversion of $x_{\text{env}}(t)$ to decibel scale, α turns from a multiplicative scale in linear space into an additive term, or offset, in logarithmic space:

$$x_{\log}(t) = 20 \cdot \log_{10} [\alpha \cdot s(t)] = 20 \cdot [\log_{10} \alpha + \log_{10} s(t)], \quad \alpha > 0 \quad (16)$$

The highpass filtering of $x_{\log}(t)$ can be approximated as a subtraction of the local signal offset within a suitable time interval $0 \ll T_{\text{HP}} < \frac{1}{f_{\text{cut}}}$:

$$x_{\text{adapt}}(t) \approx x_{\log}(t) - 20 \cdot \log_{10} \alpha = 20 \cdot \log_{10} s(t) \quad (17)$$

This eliminates α from $x_{\text{adapt}}(t)$ and thus renders it perfectly intensity-invariant, with a constant standard deviation of around 10 dB across all $\alpha > 0$ (Fig. 5c). In contrast, in the noisy case (Fig. 5b), $x_{\text{env}}(t)$ takes the form of

$$x_{\text{env}}(t) = \alpha \cdot s(t) + \eta(t), \quad x_{\text{env}}(t) > 0 \quad \forall t \in \mathbb{R} \quad (18)$$

Similar to the previous analysis (Fig. 4d), the ratio of the standard deviation of $x_{\text{env}}(t)$ to its pure-noise reference standard deviation on a double-logarithmic scale follows a constant regime for small α and a linearly increasing regime for larger α (Fig. 5d). Decibel conversion of $x_{\text{env}}(t)$

$$x_{\log}(t) = 20 \cdot \left(\log_{10} \alpha + \log_{10} \left[s(t) + \frac{\eta(t)}{\alpha} \right] \right), \quad \alpha > 0 \quad (19)$$

allows for the separation of α from $s(t)$ but introduces a scaling of $\eta(t)$ by the inverse of α , which remains present even after the offset subtraction:

$$x_{\text{adapt}}(t) \approx 20 \cdot \log_{10} \left[s(t) + \frac{\eta(t)}{\alpha} \right] \quad (20)$$

This means that, in the noisy case, α cannot be entirely eliminated from $x_{\text{adapt}}(t)$, only redistributed between $s(t)$ and $\eta(t)$. If α is sufficiently large ($\alpha \gg 1$, saturation regime), $\eta(t)$ is attenuated to the point of being negligible, so that $x_{\text{adapt}}(t)$ is a scale-free representation of $s(t)$. If α and $\eta(t)$ are at similar scales ($\alpha \approx 1$, transient regime), $x_{\text{adapt}}(t)$ largely resembles $x_{\log}(t)$. Finally, if α is sufficiently small ($0 < \alpha \ll 1$, noise regime), $\eta(t)$ masks $s(t)$ even after the intensity adaptation. Accordingly, the effective intensity invariance of $x_{\text{adapt}}(t)$ through logarithmic compression and adaptation is limited by the SNR of $x_{\text{env}}(t)$: Songs that have already sunken into the noise floor at the level of $x_{\text{env}}(t)$ cannot be recovered by subsequent processing steps. The general pattern of noise regime, transient regime, and saturation regime remains consistent across different species (Fig. 5e). However, the saturation point — the α value at which the SNR of $x_{\text{adapt}}(t)$ starts to saturate — and the saturation level — the constant SNR of $x_{\text{adapt}}(t)$ within the saturation regime — vary considerably between and within species (appendix Figs. 14+15).

For example, *C. biguttulus* and *C. mollis* display a noticeably lower saturation level compared to other species. These differences are not to be underestimated, since the saturation level of $x_{\text{adapt}}(t)$ determines the maximum input SNR for subsequent processing steps. In other words, the fact that $x_{\text{adapt}}(t)$ eventually reaches a saturation regime is, of course, desirable in the context of intensity invariance, but it also means to pass up on the higher SNR values that are achieved by $x_{\text{env}}(t)$ for the same α (up to several orders of magnitude, Fig. 5d). This trade-off between intensity invariance and SNR is a recurring phenomenon that is further addressed in the following sections.

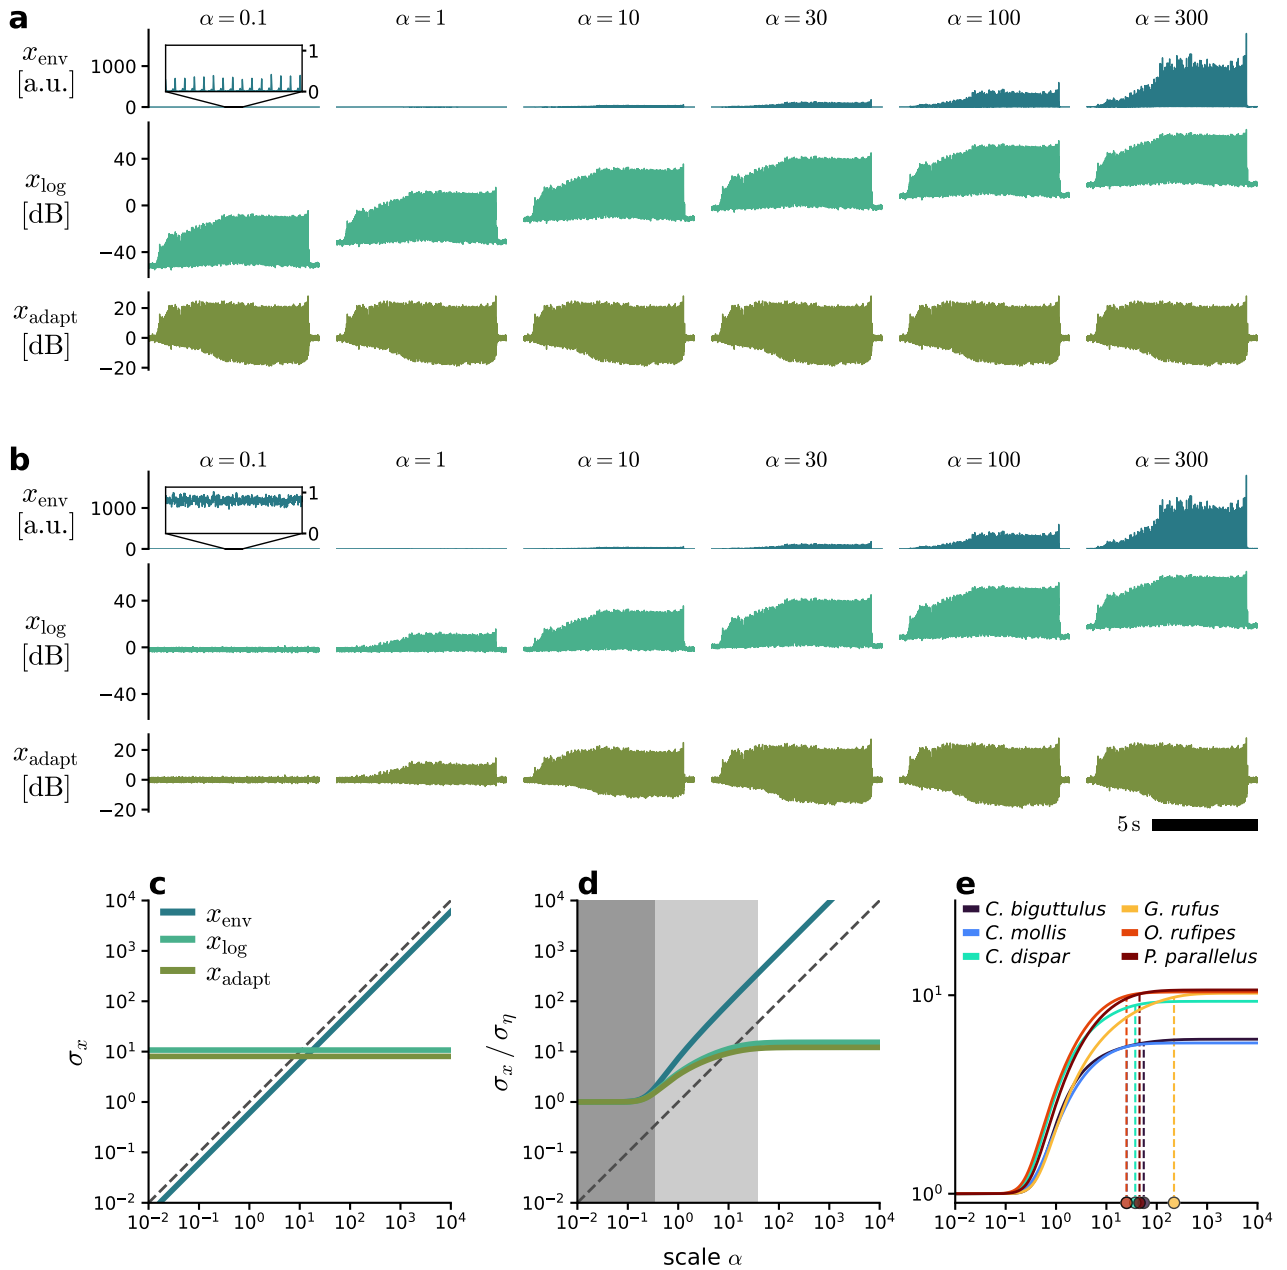


Fig. 5: Intensity invariance through logarithmic compression and adaptation is restricted by the noise floor and decreases SNR. Input $x_{\text{filt}}(t)$ consists of $s(t)$ scaled by α with optional $\eta(t)$ and is successively transformed into envelope $x_{\text{env}}(t)$, logarithmically compressed envelope $x_{\text{log}}(t)$, and intensity-adapted envelope $x_{\text{adapt}}(t)$. **Top:** Examples of $x_{\text{env}}(t)$, $x_{\text{log}}(t)$, and $x_{\text{adapt}}(t)$ for different α . **a:** Noiseless case. **b:** Noisy case. **Bottom:** Intensity measures over α . **c:** Noiseless case: Standard deviation σ_x of $x_{\text{env}}(t)$, $x_{\text{log}}(t)$, and $x_{\text{adapt}}(t)$, respectively. **d:** Noisy case: Ratio of σ_x to the respective pure-noise reference σ_η for $\alpha = 0$. Shaded areas indicate 5% (dark grey) and 95% (light grey) curve span for $x_{\text{adapt}}(t)$. **e:** Ratio of σ_x to σ_η of $x_{\text{adapt}}(t)$ as in **d** for different species (averaged over songs and recordings, appendix Fig. 14). Dots indicate 95% curve span per species.

3.1.3 Thresholding nonlinearity & temporal averaging

The third nonlinear transformation along the model pathway is the thresholding nonlinearity $H(c_i - \Theta_i)$ that transforms each kernel response $c_i(t)$ into a binary binary response $b_i(t)$,

Eq. 9. This transformation takes place after the convolutional filtering of $x_{\text{adapt}}(t)$ with kernel $k_i(t)$, Eq. 5, and is followed by the temporal averaging of $b_i(t)$ into the feature set $f_i(t)$ by a lowpass filter, Eq. 10. The effects of thresholding and temporal averaging are best illustrated based on a single kernel (Fig. 6) instead of the full set. For this analysis, input $x_{\text{adapt}}(t)$ was rescaled (Fig. 6a) and convolved with kernel $k(t)$. The resulting kernel response $c(t)$ was passed through $H(c - \Theta)$ with three different threshold values Θ (Fig. 6b-d). Each resulting binary response $b(t)$ was transformed into $f(t)$, whose average feature value μ_f serves as a measure of intensity (Fig. 6ef). The thresholding nonlinearity $H(c - \Theta)$ categorizes the values of $c(t)$ into "relevant" ($c(t) > \Theta$, $b(t) = 1$) and "irrelevant" ($c(t) \leq \Theta$, $b(t) = 0$) response values. It thereby splits the probability density $p(c, T)$ of $c(t)$ within some observed time interval T into two complementary parts around Θ :

$$\int_{\Theta}^{+\infty} p(c, T) dc = 1 - \int_{-\infty}^{\Theta} p(c, T) dc = \frac{T_1}{T}, \quad \int_{-\infty}^{+\infty} p(c, T) dc = 1 \quad (21)$$

The right-sided part of the split $p(c, T)$ corresponds to time T_1 where $c(t) > \Theta$, while the left-sided part corresponds to time $T_0 = T - T_1$ where $c(t) \leq \Theta$. The semi-definite integral over the right-sided part of $p(c, T)$ represents the ratio of T_1 relative to T because the indefinite integral of a probability density is normalized to 1. The lowpass filtering of $b(t)$ can be approximated as temporal averaging over a suitable averaging interval $T_{\text{LP}} > \frac{1}{f_{\text{cut}}}$. This allows us to express $f(t)$ as the ratio of time T_1 where $b(t) = 1$ relative to T_{LP} :

$$f(t) \approx \frac{1}{T_{\text{LP}}} \int_t^{t+T_{\text{LP}}} b(\tau) d\tau = \frac{T_1}{T_{\text{LP}}}, \quad b(t) \in \{0, 1\} \quad (22)$$

Therefore, the value of $f(t)$ at every time point t approximately signifies the cumulative probability that $c(t)$ exceeds Θ within the corresponding T_{LP} :

$$f(t) \approx \int_{\Theta}^{+\infty} p(c, T_{\text{LP}}) dc = P(c > \Theta, T_{\text{LP}}) \quad (23)$$

In a sense, $f(t)$ can be interpreted as some sort of duty cycle of $c(t)$ with respect to Θ . For example, a feature value of $f(t) = 0.4$ means that $c(t)$ exceeds Θ for approximately 40% of the time within T_{LP} . In the most extreme cases, Θ lays either above the maximum of $c(t)$ or below the minimum of $c(t)$, which results in a minimum or maximum possible feature value of $f(t) = 0$ (Fig. 6d, left column) or $f(t) = 1$, respectively.

If the time T_1 where $c(t) > \Theta$ within T_{LP} is approximately constant across t for some time interval $T_{\text{total}} > T_{\text{LP}}$, then $f(t)$ is approximately constant across $t \in T_{\text{total}}$ as well (Fig. 3c). This is fulfilled if $c(t)$ is stationary in the sense that its distribution $p(c, T_{\text{LP}})$ does not change substantially within T_{total} , which requires that T_{LP} is much longer than the relevant time scales of $c(t)$. However, stationarity of $c(t)$ is not a necessary condition for $f(t)$ to be constant because $f(t)$ depends only on the total T_1 — irrespective of the timing of individual threshold crossings — and different $p(c, T_{\text{LP}})$ can, in principle, still result in similar T_1 .

Similarly, $f(t)$ retains no information about the precise values of $c(t)$ apart from their relation to Θ . Different scales α can hence result in similar values of $f(t)$ as long as T_1 remains similar across α . The most reliable way of exploiting this invariant property of $f(t)$ is to set Θ to a value

near 0, because these values are least affected by different scales of $c(t)$. For sufficiently large α , $f(t)$ then approaches the same constant μ_f in both the noiseless and the noisy case (Fig. 6e, saturation regime). The saturation level of $f(t)$ is independent of the precise value of Θ , but the saturation point decreases with Θ (Fig. 6e). Therefore, a threshold value of $\Theta = 0$ would be the optimal choice for achieving intensity invariance at the lowest possible α . In stark contrast, the closer Θ is to 0, the higher μ_f in response to the pure noise component $\eta(t)$ and the lower the resulting SNR of $f(t)$ between noise regime and saturation regime (Fig. 6b-d, left column, and Fig. 6e). This trade-off between intensity invariance and SNR has already been observed during the previous analysis on logarithmic compression and adaptation (Fig. 5d).

Finally, the effects of thresholding and temporal averaging must be seen in the context of the previous transformation pair of logarithmic compression and adaptation: In the current analysis, the input $x_{\text{adapt}}(t)$ can be rescaled by arbitrarily large α , while in the full pathway, the current input $x_{\text{adapt}}(t)$ is the output $x_{\text{adapt}}(t)$ of the previous transformation pair and is hence capped to a maximum standard deviation of around 10 dB (Fig. 5cd). This can be illustrated by plotting μ_f not over α (Fig. 6e) but over the standard deviation of input $x_{\text{adapt}}(t)$ instead (Fig. 6f). It becomes apparent that μ_f saturates only for standard deviations of $x_{\text{adapt}}(t)$ that would already be capped. Accordingly, $f(t)$ never reaches the saturation regime as determined by the current transformation pair but rather adheres to the saturation regime determined by the previous transformation pair. In this case, the saturated μ_f is not independent of Θ anymore. The consequences of this interaction between the two mechanisms of intensity invariance are further explored in a later section.

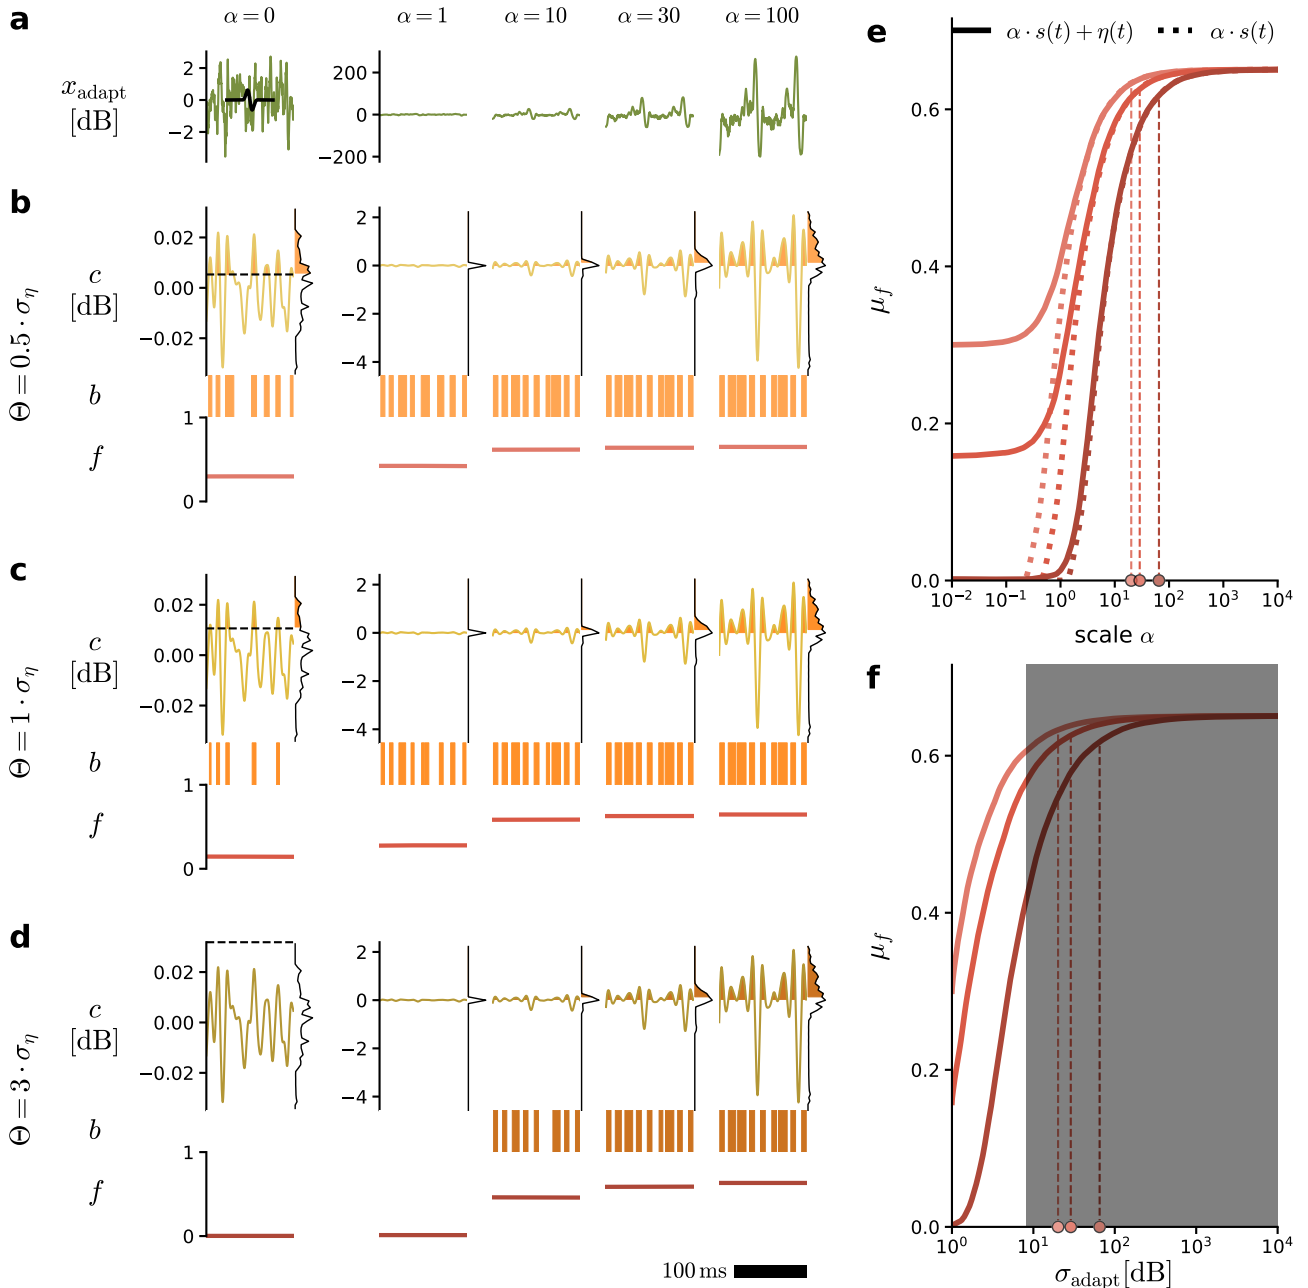


Fig. 6: Intensity invariance through thresholding and temporal averaging is mediated by the interaction of threshold value and noise floor. Input $x_{\text{adapt}}(t)$ consists $s(t)$ scaled by α with optional $\eta(t)$ and is transformed into single kernel response $c(t)$, binary response $b(t)$, and feature $f(t)$. Different color shades indicate different threshold values Θ (multiples of pure-noise standard deviation σ_η of $c(t)$ for $\alpha = 0$, with darker colors for higher Θ . See also appendix Fig. 18). **Left:** Noisy case: Examples of $x_{\text{adapt}}(t)$ as well as $c(t)$, $b(t)$, and $f(t)$ for different α . **a:** $x_{\text{adapt}}(t)$ with kernel $k(t)$ in black. **b-d:** $c(t)$, $b(t)$, and $f(t)$ based on the same $x_{\text{adapt}}(t)$ from **a** but for different Θ . **Right:** Average value μ_f of $f(t)$ for each Θ from **b-d**. Dots indicate 95% curve span (noisy case). **e:** μ_f over α , once for the noisy case (solid lines) and once for the noiseless case (dotted lines). **f:** Noisy case: μ_f over standard deviation σ_{adapt} of input x_{adapt} corresponding to α shown in **e**. Shaded area indicates values of σ_{adapt} that are capped in the output $x_{\text{adapt}}(t)$ of the previous transformation pair (Fig. 5cd).

3.2 Intensity invariance of species-specific feature representations

Having established both the meaning of the feature value and the mechanism of intensity invariance by thresholding and temporal averaging, the question remains how this mechanism acts on a set of features $f_i(t)$ based on different species-specific songs (Fig. 7a). The previous analysis was repeated with three different kernels $k_i(t)$ using a single kernel-specific threshold value Θ_i ; and the resulting average feature values μ_{f_i} were plotted over α (Fig. 7bc). Additionally, 2D feature spaces spanned by each pair of $f_i(t)$ were plotted to investigate the separability of species-specific songs based on the feature representation in dependence of α (Fig. 7de). Each species-specific combination of μ_{f_i} follows a trajectory through feature space that develops with α . These trajectories correspond to the transient regime between the constant (noise) regime and the saturation regime, which are only visible as the start and end points of the trajectories, respectively. The horizontal dashes in the colorbars indicate the range of α that corresponds to the transient regime across $f_i(t)$ for each species.

In the noiseless case, each μ_{f_i} is 0 for small α across all species (Fig. 7b) because $c_i(t)$ never exceeds Θ_i . Accordingly, each trajectory starts at the origin of the feature space (Fig. 7d). For larger α , all μ_{f_i} saturate at individual values whose combination differs between species, so that the songs of each species are eventually represented by distinct points in feature space. However, the species-specific trajectories cross each other at numerous points, which means that the songs of two species — each at a specific α — can result in the same combination of μ_{f_i} . Furthermore, the specific saturation point of $f_i(t)$ depends on the species: For *C. mollis*, all μ_{f_i} saturate around the same α , while *O. rufipes* exhibits considerable variation between the three $f_i(t)$. The larger the variation in saturation points between $f_i(t)$, the stronger the curvature of the trajectory through feature space.

In the noisy case, μ_{f_i} is non-zero even for the smallest α (Fig. 7c) because the addition of the noise component $\eta(t)$ to input $x_{\text{adapt}}(t)$ drives $c_i(t)$ above Θ_i regardless of the song component $s(t)$. The starting value of μ_{f_i} is the same across all $f_i(t)$ and species by construction of the specific Θ_i . In consequence, the trajectories through feature space do not start at the origin but rather at approximately the same point along the diagonal (Fig. 7e). For larger α , all μ_{f_i} saturate at the same values as in the noiseless case, as expected from the previous analysis (Fig. 6e). However, the trajectories now move a much shorter distance through feature space for a similar range of α due to the lower SNR of $f_i(t)$ between noise regime and saturation regime, which increases the likelihood of trajectories crossing each other. Finally, the saturation points of $f_i(t)$ for a given species are slightly higher in the noisy case, but the variation between $f_i(t)$ remains largely unchanged.

In summary, even a comparably small set of three features $f_i(t)$ can, in principle, represent different species-specific songs at distinct points in feature space, regardless of the presence of noise. However, this only holds for sufficiently large α that allow $f_i(t)$ to reach a saturation regime. During the transient regime, the species-specific combination of μ_{f_i} can very well be the same for two or more different species at specific α , although this may be alleviated by the inclusion of additional $f_i(t)$. Overall, the results of this analysis suggest that Θ_i should rather be chosen in favor of a higher SNR (Θ_i just above pure-noise $c_i(t)$) than a lower saturation point ($\Theta_i \rightarrow 0$). First, because this reduces the density of trajectories through feature space, and second, because the capping of $x_{\text{adapt}}(t)$ by the previous transformation pair likely renders

the saturation point of $f_i(t)$ less relevant.

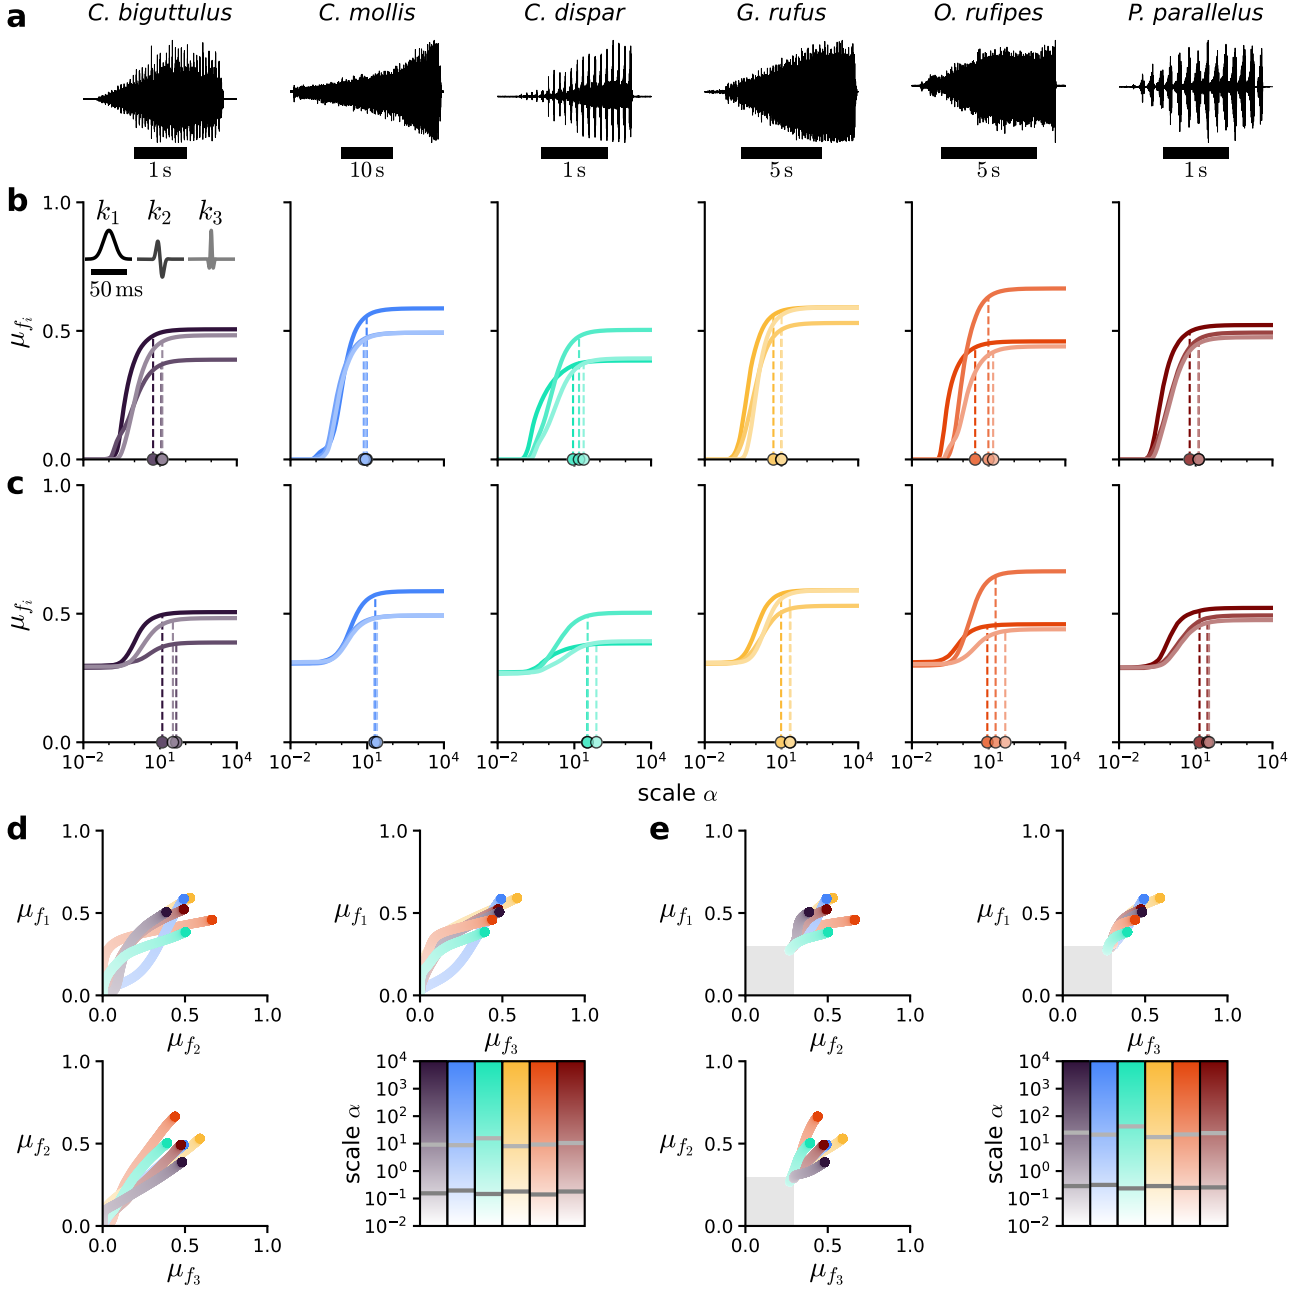


Fig. 7: Feature representation of different species-specific songs saturates at different points in feature space. Same input and processing as in Fig. 6 but with three different kernels k_i and a single kernel-specific threshold value $\Theta_i = 0.5 \cdot \sigma_{\eta_i}$ (appendix Fig. 18). **a:** Examples of species-specific grasshopper songs. **Middle:** Average value μ_{f_i} of each feature $f_i(t)$ over α per species (averaged over songs and recordings, appendix Figs. 16 and 17). Different color shades indicate different k_i . Dots indicate 95% curve span per k_i . **b:** Noiseless case. **c:** Noisy case. **Bottom:** 2D feature spaces spanned by each pair of $f_i(t)$. Each trajectory corresponds to a species-specific combination of μ_{f_i} that develops with α (colorbars). Horizontal dashes in the colorbar indicate 5% (dark grey) and 95% (light grey) curve span of the norm across all three μ_{f_i} per species. **d:** Noiseless case. **e:** Noisy case. Shaded areas indicate the average minimum μ_{f_i} across all species-specific trajectories.

3.3 Intensity invariance along the full model pathway

Through the previous analyses, we could establish two mechanisms of intensity invariance: Logarithmic compression and adaptation as well as thresholding and temporal averaging. While each transformation pair by itself can provide some level of invariance, certain results suggest that the first mechanism may actually limit or even nullify the effect of the second mechanism. In the following sections, we investigate the combined effect of both mechanisms along the full model pathway (Fig. 8) and explore the consequences of disabling the first mechanism by skipping the logarithmic compression step (Fig. 9).

3.3.1 Including logarithmic compression

For this analysis, input $x_{\text{raw}}(t)$ — including both song component $s(t)$ and noise component $\eta(t)$ — was rescaled and processed throughout all steps of the model pathway (Fig. 8a) up to the feature set $f_i(t)$. As before, the standard deviation was used as intensity measure for each resulting representation except $b_i(t)$ and $f_i(t)$. For $f_i(t)$, the average feature value μ_{f_i} was used, while $b_i(t)$ was omitted from the analysis. Plotting each intensity measure over α (Fig. 8b) reinforces many of the previous observations. For ease of visualization, the kernel-specific curves for $c_i(t)$ and $f_i(t)$ were summarized by their median. Representations prior to logarithmic compression — $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$ — show a linear increase of the intensity measure for larger α on a double-logarithmic scale. Representations after logarithmic compression — $x_{\text{log}}(t)$, $x_{\text{adapt}}(t)$, and $c_i(t)$ — are the first to reach a saturation regime and do so at approximately the same α because they are separated only by linear transformations. Feature set $f_i(t)$ reaches a saturation regime, as well. But contrary to previous results, the saturation point of $f_i(t)$ appears below that of $c_i(t)$, which suggests that the second mechanism of thresholding and temporal averaging can indeed improve intensity invariance beyond the first mechanism of logarithmic compression and adaptation. The difference in saturation points is best illustrated based on the ratio of each intensity measure to the respective pure-noise reference value (Fig. 8d). However, compressing $f_i(t)$ into a median across $k_i(t)$ conceals many kernel-specific details. It is therefore necessary to consider the development of each $f_i(t)$ over α separately (Fig. 8c). Indeed, all 40 $f_i(t)$ in the set reach a saturation regime for sufficiently large α . The saturated μ_{f_i} are distributed over a range of values — which is the prerequisite for forming species-specific combinations — but are limited to a rather small subset of possible values between 0 and 1. Based on previous results (Fig. 6f), this is likely due to the capping of $x_{\text{adapt}}(t)$ that prevents $f_i(t)$ from reaching its intrinsic saturation value; but this cannot be confirmed until the following analysis (Fig. 9). Looking at the kernel-specific SNR values of $c_i(t)$ over α (Fig. 8e) and $f_i(t)$ over α (Fig. 8f) reveals a high degree of variation between different $k_i(t)$. Certain $f_i(t)$ achieve much higher SNR values than $c_i(t)$ for the same α due to the former’s capacity for arbitrarily low pure-noise responses ($\mu_{f_i} \rightarrow 0$) and hence arbitrarily high SNR values. Finally, the question remains whether the suspected improvement of intensity invariance by $f_i(t)$ beyond $c_i(t)$ holds at the level of individual $k_i(t)$. The single saturation points based on the median across $k_i(t)$ for $c_i(t)$ and $f_i(t)$ are expanded into distributions of kernel-specific saturation points (Fig. 8g). For $c_i(t)$, the distribution is rather narrow and corresponds well to the single saturation point based on the median. For $f_i(t)$, however, the distribution is much broader and is not centered around the single saturation point based on the median but rather shifted towards lower α . Care must be taken when interpreting the height of either distribution due to the logarithmic

scaling of the underlying α axis. Nevertheless, the overall pattern suggests that the saturation points of specific $f_i(t)$ are indeed lower than those of their $c_i(t)$ counterparts. Therefore, the effect of thresholding and temporal averaging on intensity invariance is not necessarily nullified by the previous logarithmic compression and adaptation.

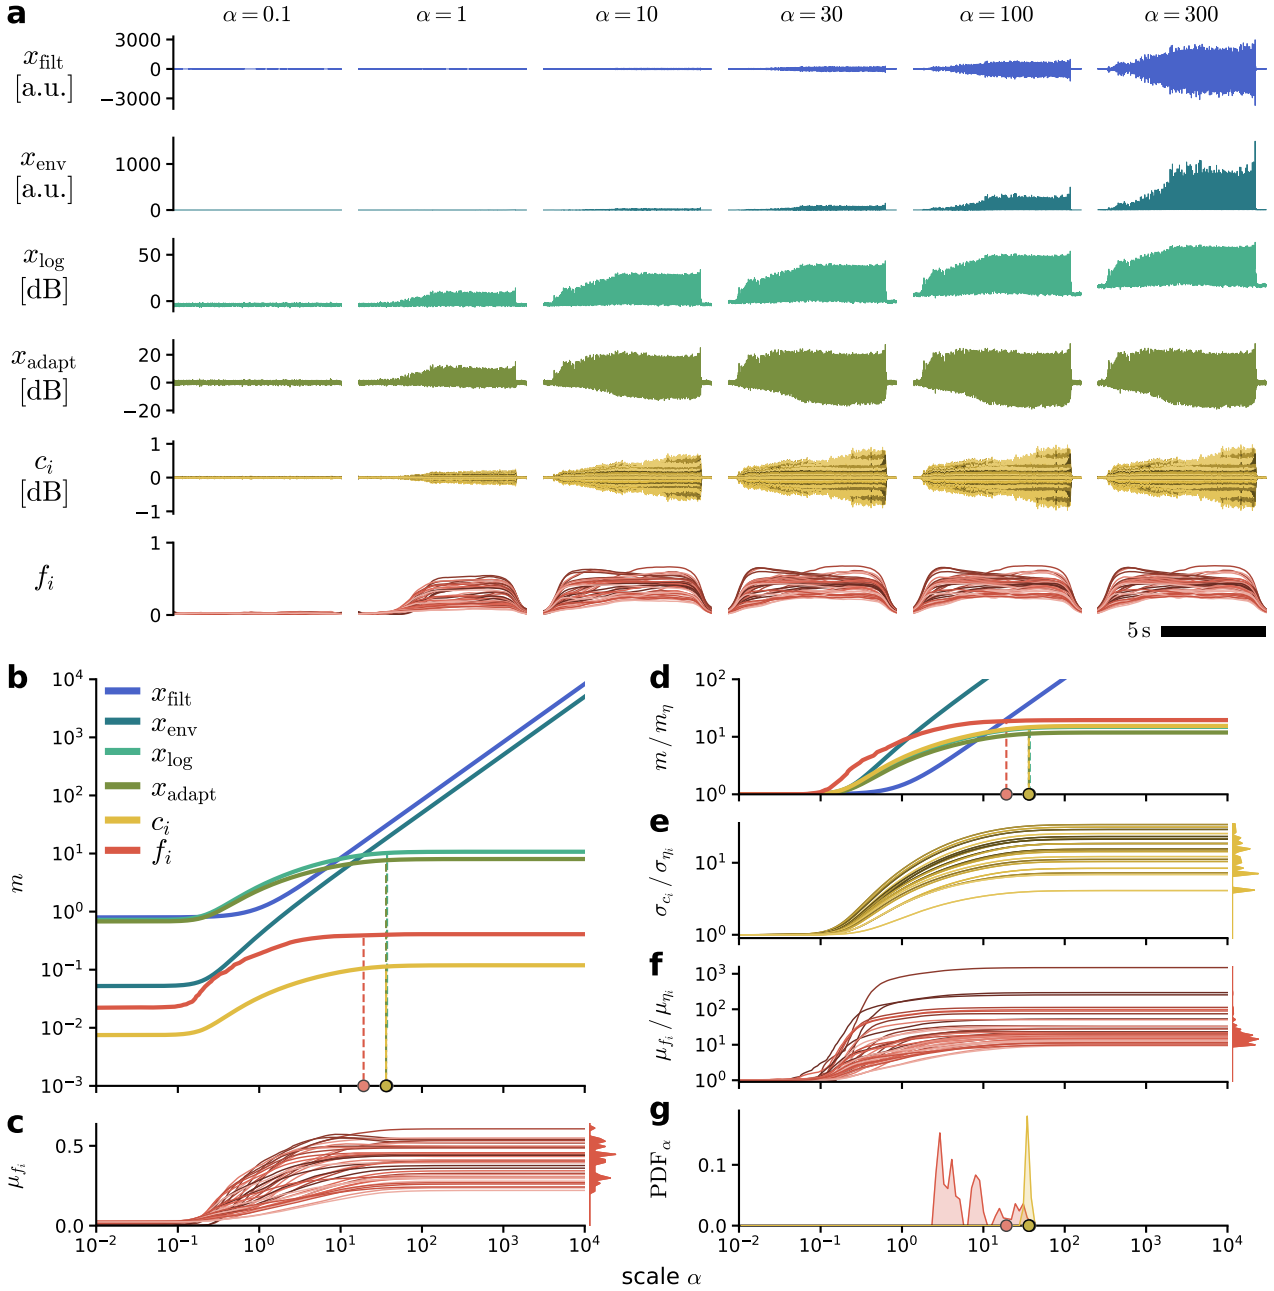


Fig. 8: Step-wise emergence of intensity-invariant song representations along the model pathway. Input $x_{\text{raw}}(t)$ consists of $s(t)$ scaled by α with added $\eta(t)$ and is processed up to the feature set $f_i(t)$ using kernel-specific threshold values $\Theta_i = 2 \cdot \sigma_{\eta_i}$ (appendix Fig.19). Different color shades indicate different types of Gabor kernels with specific lobe number n and either + or - sign, sorted (dark to light) first by increasing n and then by sign ($1 \leq n \leq 4$; first +, then - for each n ; five kernel widths σ of 1, 2, 4, 8, and 16 ms per type; 8 types, 40 $k_i(t)$ in total). **a:** Examples of $x_{\text{filt}}(t)$, $x_{\text{env}}(t)$, $x_{\text{log}}(t)$, $x_{\text{adapt}}(t)$, $c_i(t)$, and $f_i(t)$ for different α . **b:** Intensity measures over α . The median over $k_i(t)$ is shown for $c_i(t)$ and $f_i(t)$. Dots indicate 95% curve span for $x_{\text{log}}(t)$, $x_{\text{adapt}}(t)$, $c_i(t)$, and $f_i(t)$. **c:** Average value μ_{f_i} of each $f_i(t)$ over α . **d:** Ratio of intensity measures from **b** to the respective pure-noise reference for $\alpha = 0$. **e:** Ratio of standard deviation σ_{c_i} of each $c_i(t)$. **f:** Ratio of μ_{f_i} . **g:** Distributions of kernel-specific α that correspond to 95% curve span for $c_i(t)$ and $f_i(t)$. Dots indicate values based on the median from **b**.

3.3.2 Excluding logarithmic compression

The previous analysis was repeated in exactly the same way as before, except that the logarithmic compression of $x_{\text{env}}(t)$, Eq. 3, was skipped in order to disable the first mechanism of intensity invariance. Consequently, $x_{\text{adapt}}(t)$ is merely a highpass filtered version of $x_{\text{env}}(t)$; and $x_{\text{log}}(t)$ is missing entirely (Fig. 9a). As expected, all representations prior to the thresholding nonlinearity $H(c_i - \Theta_i) - x_{\text{filt}}(t)$, $x_{\text{env}}(t)$, $x_{\text{adapt}}(t)$, and $c_i(t)$ — show a linear increase of the intensity measure for larger α , while $f_i(t)$ is the only representation to reach a saturation regime (Fig. 9bd). The saturated μ_{f_i} are distributed over a much broader range of values than in the previous analysis (Fig. 9c). Intriguingly, the distribution of μ_{f_i} is symmetric around a value of 0.5. This is relevant because every kernel $k^+(t)$ in the underlying kernel set has a counterpart of opposite sign that is otherwise identical, so that $k^+(t) = -k^-(t)$. The responses of $k^+(t)$ and $k^-(t)$ to the same input $x_{\text{adapt}}(t)$ are also inverted because convolution is a linear operation: $c^+(t) = -c^-(t)$. The distributions of $c^+(t)$ and $c^-(t)$ are hence inverted to each other, as well: $p(c^+) = p(-c^-)$. Based on Eq. 23, transforming $c^+(t)$ and $c^-(t)$ further using the same Θ thus results in two complementary features $f^+(t)$ and $f^-(t)$ that are symmetric around 0.5, so that $f^+(t) = 1 - f^-(t)$. Of course, this symmetry throughout the feature representation goes hand in hand with a substantial degree of redundancy and is hardly expected to be present in the actual grasshopper auditory system. But the fact that the saturated μ_{f_i} are distributed symmetrically around 0.5 provides concrete evidence that each $f_i(t)$ is able to reach its intrinsic saturation level in the absence of logarithmic compression (Fig. 9c), which is otherwise prevented by the capping of $x_{\text{adapt}}(t)$, as seen during previous analyses (Fig. 6f and Fig. 8c). Otherwise, there appear to be no major differences in the development of $f_i(t)$ over α compared to the previous analysis, neither on the kernel-specific SNR values (Fig. 9e) nor on the distribution of kernel-specific saturation points (Fig. 9f). Overall, the most substantial consequence of skipping the logarithmic compression is that it allows $f_i(t)$ to reach its intrinsic saturation value. If this results in a wider range of μ_{f_i} across the feature set, it should be beneficial for forming species-specific combinations. However, this depends on multiple different factors such as the choice of $k_i(t)$ and Θ_i as well as the structure and distribution of the specific song and is hence not guaranteed simply by disabling logarithmic compression.

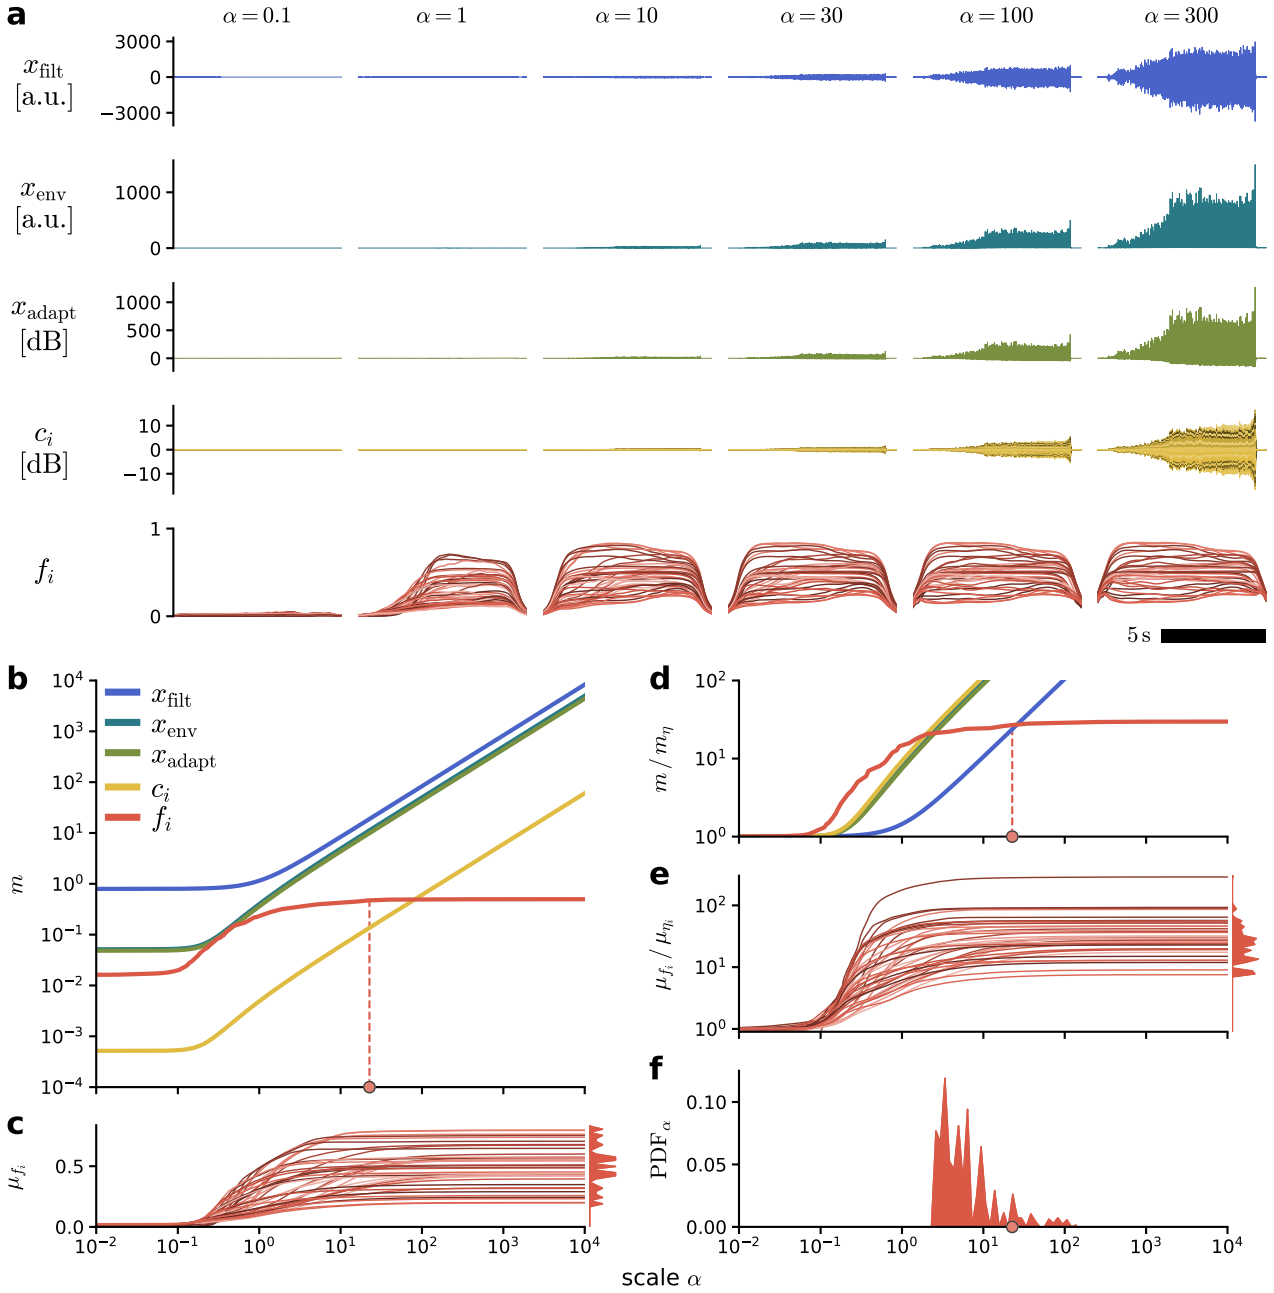


Fig. 9: Effects of disabling logarithmic compression on intensity invariance along the model pathway. Same input and processing as in Fig. 8, using kernel-specific threshold values $\Theta_i = 2 \cdot \sigma_{\eta_i}$ (appendix Fig. 20), except that logarithmic compression and hence $x_{\log}(t)$ are skipped. **a:** Examples of $x_{\text{filt}}(t)$, $x_{\text{env}}(t)$, $x_{\text{adapt}}(t)$, $c_i(t)$, and $f_i(t)$ for different α . **b:** Intensity measures over α . The median over $k_i(t)$ is shown for $c_i(t)$ and $f_i(t)$. Dot indicates 95% curve span for $f_i(t)$. **c:** Average value μ_{f_i} of each $f_i(t)$ over α . **d:** Ratio of intensity measures from **b** to the respective pure-noise reference for $\alpha = 0$. **e:** Ratio of μ_{f_i} . **f:** Distribution of kernel-specific α that correspond to 95% curve span for $f_i(t)$. Dot indicates value based on the median from **b**.

3.3.3 Intensity invariance in a naturalistic setting

So far, the analyses on intensity invariance were based on synthetically generated input signals, since these allow for a systematic manipulation of the mixture of song component $s(t)$ and noise

component $\eta(t)$ over an arbitrary range of scales α . Now, the question remains how the model pathway performs under more naturalistic conditions. The previous analysis of the full model pathway (Fig. 8) was hence repeated, using field recordings of a song of *P. parallelus* as input $x_{\text{raw}}(t)$ and a segment of background noise from the same recordings as pure-noise reference. Recordings were taken simultaneously at eight different distances d from the sender, ranging from 10 cm to 220 cm with intervals of 30 cm between microphones. The precise value of α that corresponds to a given d cannot be determined in a straightforward manner, but α is expected to be inversely proportional to d based on the inverse-square law of sound propagation. All intensity measures and ratios thereof were hence plotted over $1/d$ on a double-logarithmic scale, which is insofar comparable to previous analyses that a decade on the $1/d$ axis corresponds to a decade on the α axis. To complicate matters further, the $1/d$ axis is sampled too sparsely to determine saturation points as before based on the 95 % curve span. Instead, one has to rely on the slope of the curve to assess if, and at which $1/d$, a given representation reaches a saturation regime. Bearing these limitations in mind, the intensity measures of each representation over $1/d$ (Fig. 10b) follow a pattern that is consistent with the results of the previous simulation-based analysis (Fig. 8b): The standard deviations of $x_{\text{filt}}(t)$ and $x_{\text{env}}(t)$ increase linearly with $1/d$, respectively. The standard deviations of $x_{\text{log}}(t)$, $x_{\text{adapt}}(t)$, and $c_i(t)$ show a weaker increase with $1/d$ and appear to approach, but not reach, a saturation regime for larger $1/d$. The average feature values μ_{f_i} of $f_i(t)$ show an even weaker increase with $1/d$ and appear to reach a saturation regime for $d = 40$ cm and $d = 10$ cm, which is consistent across most $f_i(t)$ in the set (Fig. 10c). Saturation of $f_i(t)$ without saturation of $c_i(t)$ suggests that the input $x_{\text{raw}}(t)$ at the smallest $d = 10$ cm corresponds to a value of α between 10 and 20 based on comparison with the simulation-based analysis (Fig. 8b). The saturated μ_{f_i} are distributed over a comparably narrow range of values, which could in parts be a property of the songs of *P. parallelus* (see also Fig. 7bc). The ratios of each intensity measure to the respective pure-noise reference value are not aligned across representations (Fig. 10d) or kernels (Fig. 10ef) but serve to consolidate the previous observation that only $f_i(t)$ exhibits some degree of intensity invariance within the available range of $1/d$. Based on the current results, this intensity invariance of $f_i(t)$ in the field holds up to a distance of around 40 cm from the sender, decays steadily between 40 cm and 130 cm, and is substantially diminished for larger distances (Fig. 10a, bottom row).

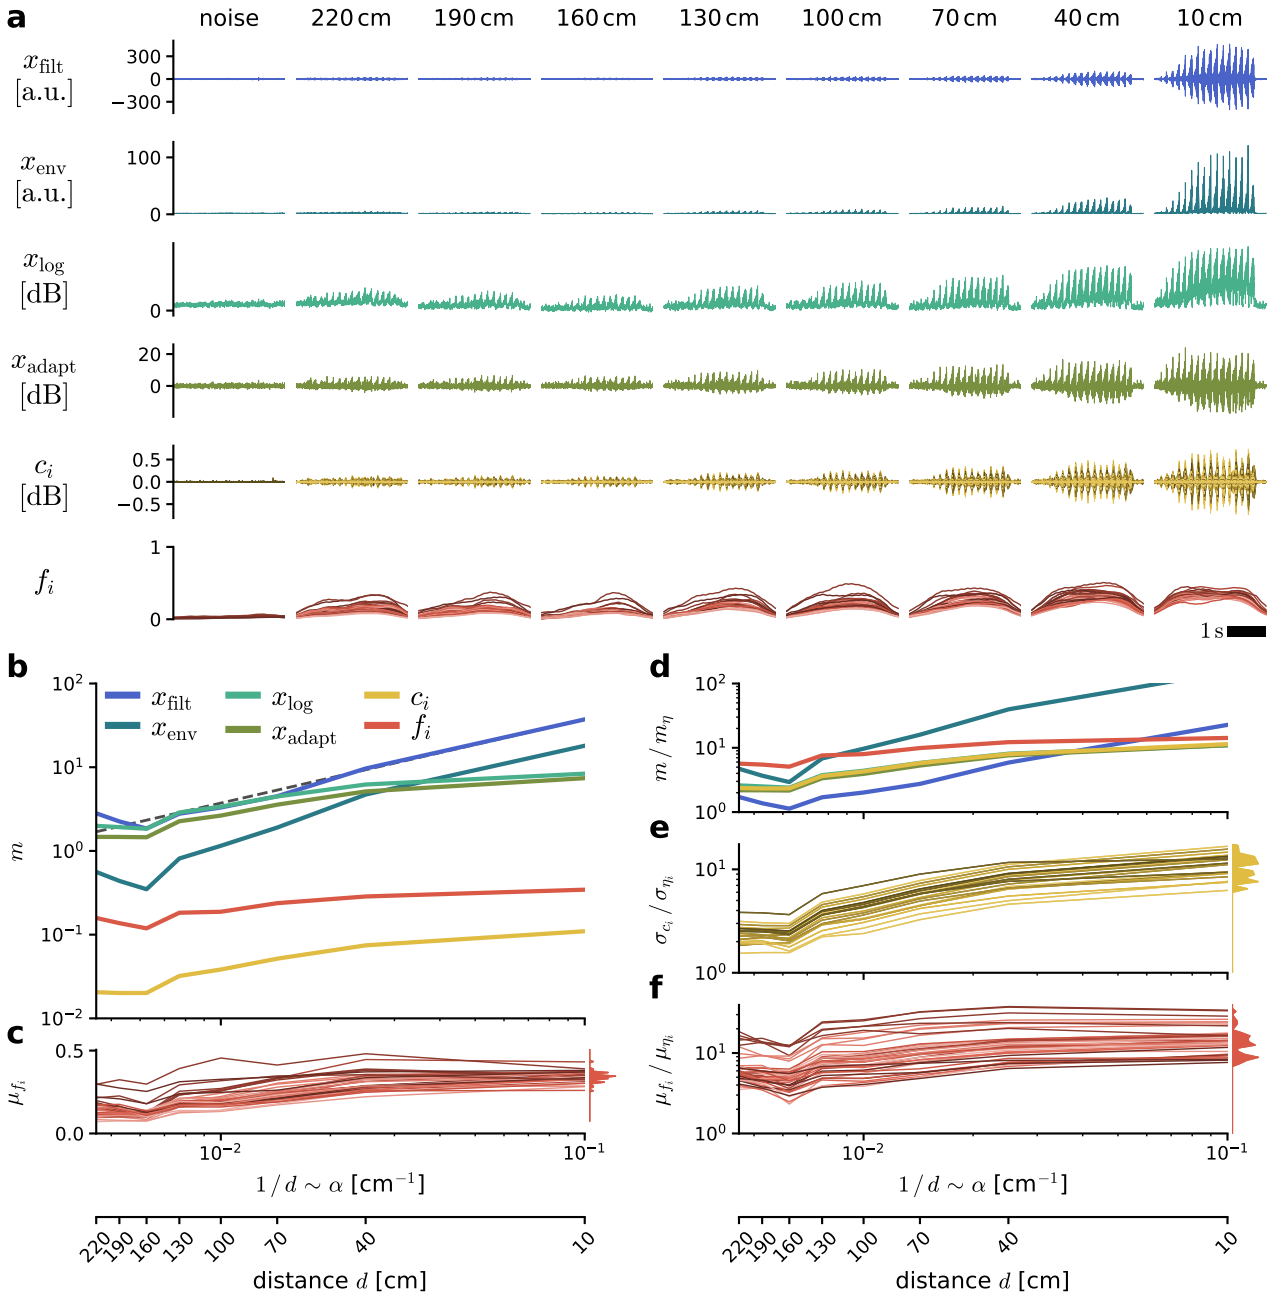


Fig. 10: Intensity invariance along the model pathway in a naturalistic setting. Input $x_{\text{raw}}(t)$ consists of a song of *P. parallelus* recorded in the field at eight different distances d and is processed up to the feature set $f_i(t)$ using kernel-specific threshold values $\Theta_i = 2 \cdot \sigma_{\eta_i}$ (appendix Fig. 21). Different color shades indicate different types of Gabor kernels with specific lobe number n and either + or - sign, sorted (dark to light) first by increasing n and then by sign ($1 \leq n \leq 4$; first +, then - for each n ; five kernel widths σ of 1, 2, 4, 8, and 16 ms per type; 8 types, 40 $k_i(t)$ in total). **a:** $x_{\text{filt}}(t)$, $x_{\text{env}}(t)$, $x_{\text{log}}(t)$, $x_{\text{adapt}}(t)$, $c_i(t)$, and $f_i(t)$ at each d . A noise segment from the same recording is shown for reference. **b:** Intensity measures over d . The median over $k_i(t)$ is shown for $c_i(t)$ and $f_i(t)$. **c:** Average value μ_{f_i} of each $f_i(t)$ over d . **d:** Ratio of intensity measures from **b** to the respective value obtained from the noise reference. **e:** Ratio of standard deviation σ_{c_i} of each $c_i(t)$. **f:** Ratios of μ_{f_i} .

3.4 Interspecific and intraspecific feature variability

In the final analysis of the current study, we investigated the variability of songs in the feature representation between different species and within the same species (Fig. 11). Naturally, a feature representation that is both consistent across different songs of the same species and sufficiently different between songs of different species is a fundamental prerequisite for species-specific song recognition. The data used in this analysis corresponds to the saturated μ_{f_i} of each $f_i(t)$ from the previous analysis of the full model pathway (Fig. 8c), using different songs of *O. rufipes* for the intraspecific comparisons and single songs from a number of species for the interspecific comparisons (also shown in Fig. 7a). Accordingly, each song is represented by 40 values of μ_{f_i} based on the same set of $f_i(t)$. For each comparison, μ_{f_i} from one song was plotted against μ_{f_i} from the other song, so that each dot within a subplot corresponds to a single feature $f_i(t)$. For the intraspecific comparisons (Fig. 11, upper triangular), the pairs of μ_{f_i} are distributed closely around the diagonal, with a minimum correlation coefficient of $\rho = 0.82$, a maximum of $\rho = 0.99$, and a median of $\rho = 0.91$. A given $f_i(t)$ thus tends to have a similar μ_{f_i} across different songs of the same species. In contrast, the pairs of μ_{f_i} for the interspecific comparisons (Fig. 11, lower triangular) are distributed in a variety of different ways, most in broader clouds (e.g. *C. biguttulus* vs. *C. mollis*) but some more narrowly around the diagonal (e.g. *P. parallelus* vs. *C. dispar*). The correlation coefficients ρ vary widely between different interspecific comparisons, with a minimum of $\rho = -0.1$, a maximum of $\rho = 0.91$, and a median of $\rho = 0.40$. A given $f_i(t)$ therefore tends to have a less similar μ_{f_i} across different species than within the same species, although certain exceptions exist (Fig. 11, lower right). Accordingly, the feature representation that is generated by the model pathway is, in principle, suitable for the distinction between different species-specific songs. However, even the songs of the same species are subject to considerable variability in various aspects and depending on a multitude of external and internal factors, which cannot be fully captured based on a limited number of songs. The results of the current analysis are hence to be treated as a proof-of-concept that paves the way towards more comprehensive investigations on the details of song representation in feature space, including the effects of different parameters of the model pathway as well as the inclusion of additional songs and species to reflect the complexity of natural song variation.

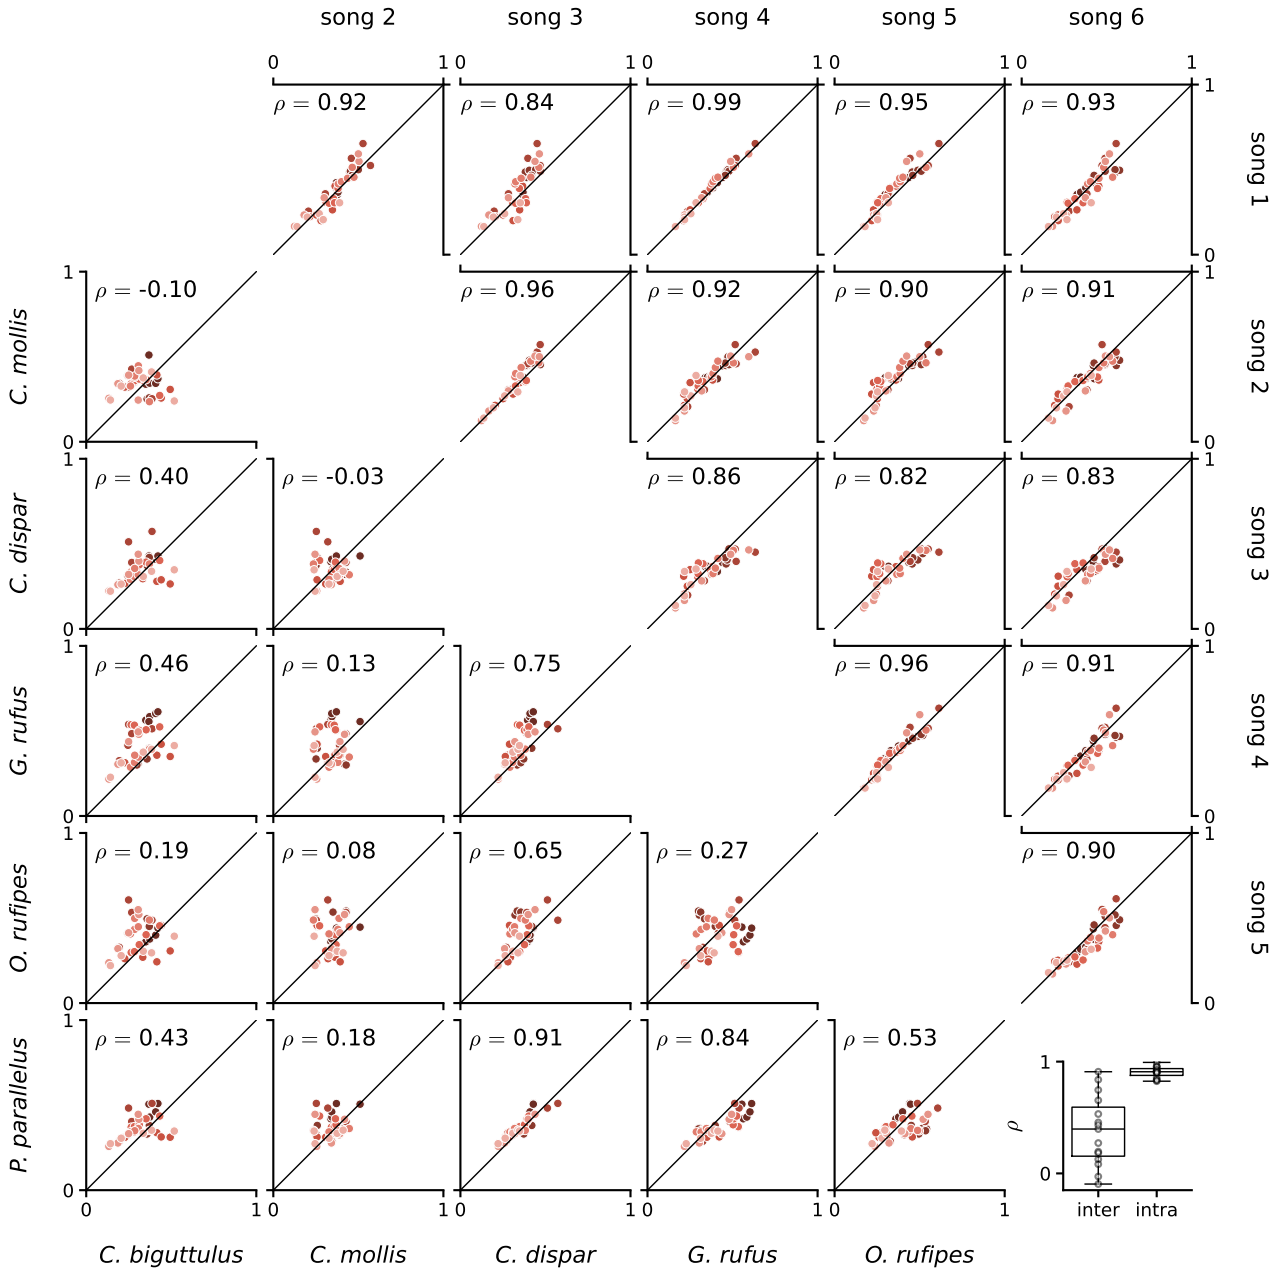


Fig. 11: Interspecific and intraspecific feature variability. Average value μ_{f_i} of each feature $f_i(t)$ against its counterpart from a 2nd feature set based on a different input $x_{\text{raw}}(t)$. Data is based on the saturated μ_{f_i} from Fig. 8. Each dot within a subplot represents a single $f_i(t)$. Different color shades indicate different types of Gabor kernels with specific lobe number n and either + or - sign, sorted (dark to light) first by increasing n and then by sign ($1 \leq n \leq 4$; first +, then - for each n ; five kernel widths σ of 1, 2, 4, 8, and 16 ms per type; 8 types, 40 kernels in total). **Lower triangular:** Interspecific comparisons between single songs of different species. **Upper triangular:** Intraspecific comparisons between different songs of a single species (*O. rufipes*). **Lower right:** Distribution of correlation coefficients ρ for each interspecific and intraspecific comparison. Dots indicate single ρ values.

4 Discussion

In the current study, we have established a physiologically inspired functional model of the grasshopper song recognition pathway; from the sound reception at the tympanal membrane over peripheral receptor neurons and local interneurons to the generation of a high-dimensional feature representation at the level of the ascending neurons and beyond in the SEG. Using this model pathway, we have identified two computational key mechanisms for the emergence of intensity-invariant song representations. Each mechanism comprises a nonlinear transformation and a subsequent linear transformation. The first mechanism consists of logarithmic compression and adaptation by highpass filtering, which takes place at the level of the receptor neurons and local interneurons. The second mechanism consists of thresholding and temporal averaging by lowpass filtering, which takes place either at the level of the ascending neurons or further downstream in the SEG. Systematic investigation of both mechanisms revealed a persistent trade-off between the intensity invariance and the SNR of the song representations along the pathway. In the following, we briefly reflect on the potential of functional modelling for research on sensory systems. We then discuss the implications of our findings for the evolutionary design of both the auditory system and the species-specific songs of grasshoppers as well as the ethological relevance of intensity invariance in a natural acoustic environment.

4.1 Leveraging functional modelling to investigate sensory systems

Our understanding of sensory processing systems is based on the distributed accumulation of anatomical, physiological, and ethological evidence. Functional modelling provides a powerful tool to integrate the available fragments into a coherent whole. It facilitates systematic, reproducible investigations of relevant parameters, such as scale α or threshold value Θ_i . It also allows to address questions of broader scope by generalizing from concrete evidence. For instance, the interaction between the two mechanisms of intensity invariance is most assessible if both mechanisms can be treated as consecutive stages along the pathway — where the output of the first stage relates directly to the input of the second stage — rather than separate entities. Moreover, the model pathway provides a general basis for comparing song representations across different species without the need for species-specific models. However, the potential of functional modelling for research on a sensory system depends entirely on the amount of available knowledge about the system and its specific stimuli. The grasshopper song recognition pathway is a comparably simple, extensively researched and hence well-understood sensory system and is therefore a particularly suitable candidate for functional modelling. Other sensory systems that are either more complex or have not been subject to decades of study will likely not be suitable for this approach yet.

4.2 Repetitive song structure and temporal averaging as design principle for a robust feature representation

The songs of different species are represented by specific combinations of feature values, which should be as constant as possible for the duration of a song to facilitate recognition. Feature $f_i(t)$ is constant if the time where kernel response $c_i(t)$ exceeds the threshold value Θ_i within the averaging interval T_{LP} is the same at each time point $t \in T_{total} > T_{LP}$. This is fulfilled if $c_i(t)$ is stationary, so that its distribution $p(c_i, T_{LP})$ does not change substantially within T_{total} .

Grasshoppers sing by pulling the stridulatory file on the hindlegs across a resonating vein on the forewings (Otto von Helversen and Elsner 1977; Stumpner and Otto von Helversen 1994; D. v. Helversen and Otto von Helversen 1997). Different stridulatory motor patterns allow for the production of vastly different song patterns without modifying the overall stridulation apparatus (Stumpner and Otto von Helversen 1994). The song pattern could hence be changed frequently and substantially throughout the song; yet many species resort to songs with a regular, highly repetitive structure (Tishechkin and Bukhvalova 2009). From the perspective of the model pathway, a repetitive song pattern is necessary because it translates into a repetitive structure of $c_i(t)$. If the song pattern is sufficiently regular, $c_i(t)$ is assumed to be stationary, which lays the basis for constant $f_i(t)$ and hence reliable recognition throughout the song. This is particularly relevant for the calling songs — whose primary function is the broadcasting of species identity — whereas the courtship songs tend towards a slightly more complex structure (V Yu Vedenina and Helversen 2003; Varvara Yu Vedenina and Mague 2011; V Yu Vedenina and Shevchenko 2012). Different accounts of how the ancestral calling song could have looked like agree that it likely possessed a repetitive structure (Von Helversen and Von Helversen 1994; Varvara Yu Vedenina and Mague 2011; Mague and Sevastianov et al. 2023). This suggests that the song recognition pathway has long been evolving around repetitive song patterns. The calling songs of many extant species — while extremely diverse in their details (Tishechkin and Bukhvalova 2009) — might still have to conform to this design principle in order to be recognizable. However, there are exceptions: For example, *Chorthippus dorsatus* produces a composite calling song that consists of a repetitive syllable-pause pattern followed by a noisy rattling sound (Stumpner and Otto von Helversen 1992; Stumpner and Otto von Helversen 1994). Females respond to the song only if both parts are present (Stumpner and Otto von Helversen 1992). It has been suggested that the second part of the song instead serves as a fitness signal, so that females might recognize the song based on the first part but choose not to respond if the second part is missing.

The minimum T_{LP} must encompass at least a few cycles of $c_i(t)$ to ensure a stationary $p(c_i, T_{LP})$. Behavioral experiments have shown that recognition of the song pattern is not drastically impaired if the duration of some syllables is changed (D. v. Helversen and Otto von Helversen 1998). In particular, recognition was least impaired if the average replacement duration corresponded roughly to the original duration, even though the individual replacements were much shorter or longer. Accordingly, the more cycles of $c_i(t)$ are included in T_{LP} , the more robust $f_i(t)$ is against irregularities in the song pattern. However, the longer T_{LP} , the longer $f_i(t)$ takes to stabilize after the onset of the song due to the inclusion of noise, which narrows the time window during which $f_i(t)$ is constant. If T_{LP} exceeds the duration of the song, $f_i(t)$ will never be constant at all. The optimal T_{LP} for a specific song is hence determined by the duration of a typical syllable-pause cycle (lower bound) and the total song duration (upper bound). Both parameters vary widely across different grasshopper species (Tishechkin and Bukhvalova 2009), which suggests that the optimal T_{LP} is likely species-specific. In the model pathway, T_{LP} is fixed at around 1 second ($f_{cut} = 1$ Hz), so that $f_i(t)$ takes accordingly long to stabilize. In contrast, *C. biguttulus* has been shown to respond to songs that consist of only 3 syllable-pause cycles and are merely 250 ms long (Ronacher and Krahe 1998), which suggests a much shorter T_{LP} in this species. It appears plausible that grasshoppers recognize conspecific songs not by a singular combination of feature values (a point in feature space) but within a certain tolerance (a region in feature space). Song responsiveness in grasshoppers is subject to a speed-accuracy trade-off (Clemens, Bernhard Ronacher, and Reichert 2021) — a grasshopper could thus either

respond as soon as $f_i(t)$ is within tolerance or wait for $f_i(t)$ to stabilize for additional certainty.

4.3 Invariant processing in the grasshopper auditory system

The notion of invariance is fundamental for sensory processing systems. Invariance, in the general sense, can be described as the property of a transformation to maintain variation across certain meaningful input parameters in its output while discarding variation across other input parameters. This boils down to a selective input-output decorrelation that allows the system to represent only those aspects of the stimulus that are behaviorally relevant to the organism.

The grasshopper auditory system has to deal with a number of sources of non-informative song variation. For instance, the temporal structure of the song pattern warps with temperature (Skovmand and Boel Pedersen 1983). The auditory system can compensate for this time warping by reading out relative temporal relationships, such as the ratio of syllable duration to pause duration, rather than the absolute time intervals (Creutzig, Wohlgemuth, et al. 2009; Creutzig, Benda, et al. 2010). This allows for reliable song recognition across different temperatures (D. v. Helversen 1972). Accordingly, the auditory system does likely not retain any information about the precise duration of syllables and pauses.

The situation is more complex for variations in song intensity. Song intensity at the receiver’s position depends mostly on the distance to the sender and is therefore not a reliable cue to infer species identity. The auditory system must hence be invariant to intensity variations to recognize conspecific songs regardless of sender distance. However, song intensity — specifically, the interaural intensity difference — is also a relevant cue for directional hearing, which is essential for phonotaxis (D. v. Helversen and Rheinlaender 1988). Interference between song recognition and directional hearing is avoided in the auditory system by distributing both functions across two parallel pathways (D. v. Helversen 1984; Bernhard Ronacher, D. v. Helversen, and Helversen 1986). This is the main reason why our model pathway is focused entirely on song recognition and has no capacity for directional hearing, even though it is crucial to the grasshopper’s behavior.

Song intensity is a function of the song amplitudes within a certain time frame. ”Invariance to variations in song intensity” is hence entirely a matter of time scales. It can refer to intensity variations across different songs (longer time scales) or intensity variations across the syllables within a song (shorter time scales), but also to the intensity difference that differentiates a syllable from a pause (very short time scales). The time scale of intensity invariance must therefore be sufficiently long to leave the syllables and pauses of the song pattern intact. In the model pathway, this time scale-selectivity is reflected by the cutoff frequency f_{cut} of the highpass filter that underlies the adaptation of $x_{\text{adapt}}(t)$: Most f_{cut} except the lowest ones are effective in removing the local offset of $x_{\text{log}}(t)$ and render $x_{\text{adapt}}(t)$ intensity-invariant, but only sufficiently low f_{cut} preserve the relevant amplitude dynamics of the song pattern. The time scale of intensity invariance by thresholding and temporal averaging is determined by the averaging interval T_{LP} . However, unlike f_{cut} , T_{LP} is not constrained by the need to preserve the song pattern but rather to provide a suitable degree of temporal integration (Section 4.2).

4.4 Intensity invariance versus SNR along the model pathway

A stridulating grasshopper generates a song with a specific initial intensity, which is steadily attenuated as the song propagates through the environment (Michelsen 1978). A listening grasshopper receives a sound signal $x_{\text{raw}}(t)$, which is a mixture of the song component $s(t)$ with scale α and the environmental noise component $\eta(t)$. The greater the distance between sender and receiver, the smaller α and hence the lower the SNR of $x_{\text{raw}}(t)$ at the position of the receiver. The tympanal bandpass filtering of $x_{\text{raw}}(t)$ into $x_{\text{filt}}(t)$ likely improves the SNR by attenuating frequencies outside the relevant range of grasshopper songs. The SNR is further improved by the rectification and lowpass filtering of $x_{\text{filt}}(t)$ into $x_{\text{env}}(t)$. The lower the cutoff frequency f_{cut} of the lowpass filter, the higher the SNR of $x_{\text{env}}(t)$ for a given α , although f_{cut} must also be sufficiently high to preserve the amplitude dynamics of the song pattern. The first processing steps along the pathway are hence designed to improve the SNR of the song representation beyond the initial SNR of $x_{\text{raw}}(t)$.

The first mechanism of intensity invariance consists of logarithmic compression and adaptation of $x_{\text{env}}(t)$ into $x_{\text{adapt}}(t)$. In the absence of $\eta(t)$, $x_{\text{adapt}}(t)$ is a perfectly intensity-invariant representation of $s(t)$. In the presence of $\eta(t)$, $x_{\text{adapt}}(t)$ is intensity-invariant only for a sufficiently high SNR of $x_{\text{env}}(t)$. The preceding SNR improvements from $x_{\text{raw}}(t)$ to $x_{\text{env}}(t)$ thus serve to improve the intensity invariance of $x_{\text{adapt}}(t)$ by shifting the saturation point towards lower α . However, this effect is limited — if the SNR of $x_{\text{raw}}(t)$ at the receiver’s position does not allow for a sufficiently high SNR of $x_{\text{env}}(t)$, $x_{\text{adapt}}(t)$ will not be intensity-invariant. In this case, the receiver is presumably less likely to recognize $x_{\text{raw}}(t)$ as a conspecific song. The limitation of the intensity invariance of $x_{\text{adapt}}(t)$ by the SNR of $x_{\text{raw}}(t)$ might hence at least in parts be responsible for the limited maximum distance at which song recognition is possible (Lang 2000) and the selection towards song patterns that are robust to noise masking (Einhäupl et al. 2011).

The SNR of each song representation prior to $x_{\text{adapt}}(t)$ increases monotonically with α (excluding $0 < \alpha \ll 1$, noise regime). These representations maintain the full extent of the initial SNR of $x_{\text{raw}}(t)$ and hence never achieve intensity invariance. In contrast, the SNR of $x_{\text{adapt}}(t)$ saturates for sufficiently high α . Accordingly, $x_{\text{adapt}}(t)$ is intensity-invariant but cannot have a higher SNR than the saturation level, which indicates a fundamental trade-off. The saturation level of $x_{\text{adapt}}(t)$ varies across species and songs. This variability is likely rooted in the way in which logarithmic compression acts on the specific $x_{\text{env}}(t)$, which depends on the f_{cut} of the lowpass filter as well as the temporal structure and frequency spectrum of the rectified $x_{\text{filt}}(t)$. Across all songs, the saturation level of $x_{\text{adapt}}(t)$ has never been observed to exceed a SNR of around 10. This is a substantial reduction in SNR, considering that the SNR of preceding representations had been orders of magnitude higher. Part of this reduction stems from the amplification of smaller values of $x_{\text{env}}(t)$ by logarithmic compression, which raises the noise floor of $x_{\text{adapt}}(t)$ relative to the song. Accordingly, the low SNR of $x_{\text{adapt}}(t)$ appears to be a necessary price to pay for its intensity invariance through logarithmic compression and adaptation. But the trade-off between intensity invariance and SNR likely goes beyond the particular mechanisms along the pathway. After all, a transformation is not expected to compress a range of different input intensities into a constant output intensity without sacrificing some of the corresponding input SNR. This suggests that the trade-off is a more general principle that applies to any transformation that achieves or improves intensity invariance.

The second mechanism of intensity invariance consists of thresholding and temporal averaging of $c_i(t)$ into $f_i(t)$. Here, the trade-off between intensity invariance and SNR is mediated by the threshold value Θ_i . A lower Θ_i ($\Theta_i \rightarrow 0$) improves the intensity invariance of $f_i(t)$ by shifting the saturation point towards lower α . However, a lower Θ_i also raises the noise floor of $f_i(t)$ by including more of the pure-noise $c_i(t)$, which decreases the SNR of $f_i(t)$. The distribution $p(c_i, T_{LP})$ of the pure-noise $c_i(t)$ is very close to a normal distribution with mean $\mu \approx 0$ for all kernels $k_i(t)$. The value of the pure-noise $f_i(t)$ is hence 0.5 for $\Theta_i = 0$ and decreases for higher Θ_i . If Θ_i is set above the maximum of $c_i(t)$, the pure-noise feature value is 0, which results in an "unlimited" SNR of $f_i(t)$. In this case, any non-zero feature value that is sustained for a sufficient duration could serve as indicator for the presence of $s(t)$ in $x_{\text{raw}}(t)$, although at the cost of a higher saturation point. Of course, this would require a fine evolutionary tuning of Θ_i to the properties of the natural noise in a certain habitat to avoid false positives.

The saturation level of $f_i(t)$ is independent of Θ_i as long as the intensity invariance by the previous mechanism is neglected.

If $f_i(t)$ can achieve an arbitrarily high SNR, it can counteract the comparably low SNR of $x_{\text{adapt}}(t)$

The maximum of the pure-noise $c_i(t)$ is assumed to be very small due to the various SNR improvements along the pathway, so that the required increase in Θ_i and hence the saturation point of $f_i(t)$ is not expected to be substantial.

Two consecutive mechanisms of intensity invariance do not necessarily add up to a stronger overall intensity invariance. If the first mechanism results in a lower saturation point than the second mechanism by itself, the saturation point of feature $f_i(t)$ will be determined solely by the first mechanism. In this case, the saturation level of $f_i(t)$ will conform to the intensity that $f_i(t)$ can reach for the given saturation point rather than the intrinsic saturation level of $f_i(t)$. Conversely, if the second mechanism results in a lower saturation point than the first mechanism, both the saturation point and the saturation level of $f_i(t)$ will be determined by the second mechanism. The saturation points of $f_i(t)$ across the set are distributed over a much wider range than those of the preceding kernel responses $c_i(t)$, which suggests that the interaction between the two mechanisms is specific to individual kernels. A number of $f_i(t)$ achieve a lower saturation point than the respective $c_i(t)$, whereas some $f_i(t)$ exhibit similar or only marginally lower saturation points. In these cases, the question arises to what extent two consecutive mechanisms of intensity invariance are actually beneficial for the overall system.

From a computational perspective, the answer could be that logarithmic compression and adaptation is a necessary preprocessing step towards robust $f_i(t)$ because it works towards a more consistent distribution $p(c_i, T_{LP})$ of $c_i(t)$. If $p(c_i, T_{LP})$ is consistent between different songs of the same species, a static threshold value Θ_i is sufficient to generate a consistent species-specific feature representation. If $p(c_i, T_{LP})$ is consistent over the course of a song, $f_i(t)$ is constant throughout the song, which extends the time window for reliable recognition (Section 4.2).

First, the preprocessing results in a more consistent distribution $p(c_i, T_{LP})$ of $c_i(t)$ between songs of different intensity and in turn allows for the generation of consistent $f_i(t)$ under a static threshold value Θ_i . Second, this preprocessing improves the temporal regularity of the song pattern by mitigating the slow "ramping" amplitude modulation that is common to many

grasshopper songs.

This preprocessing likely improves the temporal regularity of the song pattern in $x_{\text{adapt}}(t)$ and $c_i(t)$, which is required for constant $f_i(t)$ across the duration of a song (Section 4.2).

From a physiological perspective, the answer is likely that neurons possess only a limited firing rate for encoding stimulus intensities that can range over several orders of magnitude. Sigmoidal tuning curves over logarithmically compressed stimulus intensities are a common property of sensory neurons across various modalities (SOURCE?), and neurons of the grasshopper auditory system are no exception (Suga 1960; Gollisch, Schütze, et al. 2002).

→ Ronacher and R. Hennig 2004

4.5 Implications for behavior in a natural acoustic environment

Most grasshoppers live in environments that are communally inhabited by numerous individuals from multiple species. Their acoustic environment is characterized by noise from various sources — abiotic ones like wind and water, but also the songs of both hetero- and conspecifics. This limits the SNR that each individual can achieve for its own song, and hence the effectiveness of the intensity-invariant processing in the auditory system. Producing higher song intensities is not a viable solution to this problem, because these also contribute to the overall noise floor. A possible behavioral solution could be to produce songs in a "turn-taking" manner to avoid the temporal superposition of multiple songs into overly intense signals. This would also prevent the mutual distortion of the respective song pattern. Another solution could be to spatially separate from other nearby grasshoppers to spread the potential noise sources over a larger area. However, according to our analysis based on field recordings as well as previous work on the topic (Lang 2000), reliable song recognition is limited to little more than 1 m from the sender, so that a grasshopper also cannot afford to stay too far away from its conspecifics. A better solution may hence be to collectively produce songs at lower-than-possible intensities, which would reduce the overall noise floor for all nearby individuals. Importantly, the limitation of intensity invariance by SNR likely applies to all grasshoppers regardless of species, so that the behavioral strategies could be shared among the species that coexist in a given habitat.

Stange and Ronacher 2012?

Krämer and Ronacher 2018

Einhäupl et al. 2011

Snedden et al. 1998?

Tishechkin and Bukhvalova 2009?

5 Appendix




Fig. 12: Conversion of the noise component by envelope extraction. Standard deviation σ_η of noise component $\eta(t)$ within the signal envelope $x_{\text{env}}(t)$ over scale α . Based on input $x_{\text{raw}}(t)$ with $\sigma_\eta = 1$ (corresponding to the analysis underlying Fig. 4), using 100 random realizations of $\eta(t)$.




Fig. 13: Species-specific data underlying Fig. 4e. Ratio of the standard deviation σ_{env} to the pure-noise reference σ_η of the signal envelope $x_{\text{env}}(t)$ over scale α for different cutoff frequencies f_{cut} of the lowpass filter extracting $x_{\text{env}}(t)$. Solid lines and shaded areas indicate mean \pm standard deviation across songs per recording. Dashed lines indicate mean across recordings (shown in Fig. 4e).

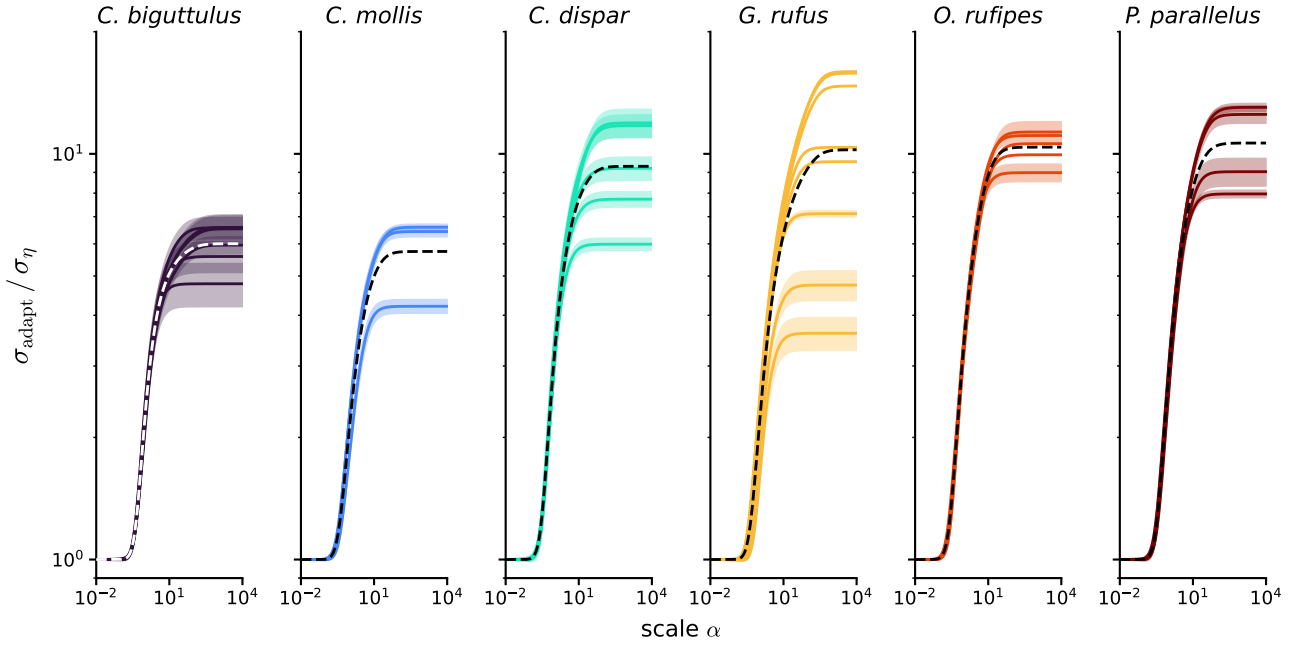


Fig. 14: Species-specific data underlying Fig. 5e. Ratio of the standard deviation σ_{adapt} to the pure-noise reference σ_{η} of the intensity-adapted envelope $x_{\text{adapt}}(t)$ over scale α . Solid lines and shaded areas indicate mean \pm standard deviation across songs per recording. Dashed lines indicate mean across recordings (shown in Fig. 5e).

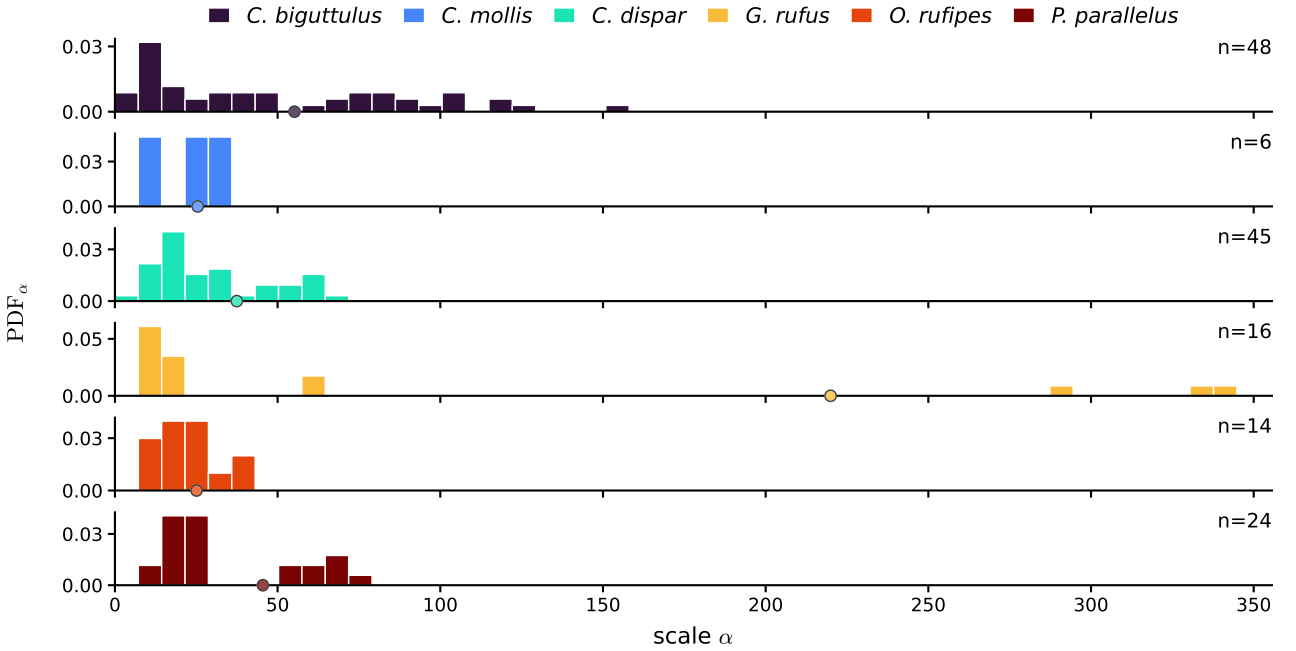


Fig. 15: Species-specific saturation points underlying Fig. 5e. Distribution of saturation points (95% curve span) of ratio $\sigma_{\text{adapt}}/\sigma_{\eta}$ of the intensity-adapted envelope $x_{\text{adapt}}(t)$ over scale α across all available songs. Dots indicate the saturation point of the mean curve across songs and recordings (shown in Fig. 5e, see also appendix Fig. 14).




Fig. 16: Species-specific data underlying Fig. 7bd. Average value μ_{f_i} of each of the three features $f_i(t)$ over scale α in the noiseless case. Solid lines and shaded areas indicate mean \pm standard deviation across songs per recording. Dashed lines indicate mean across recordings (shown in Fig. 7bd).




Fig. 17: Species-specific data underlying Fig. 7ce. Average value μ_{f_i} of each of the three features $f_i(t)$ over scale α in the noisy case. Solid lines and shaded areas indicate mean \pm standard deviation across songs per recording. Dashed lines indicate mean across recordings (shown in Fig. 7ce).




Fig. 18: Relation between threshold value and pure-noise feature value for Fig. 6 and Fig. 7. Proportion of pure-noise kernel response $c_i(t)$ that exceeds threshold value Θ_i — which determines the average value μ_{f_i} of feature $f_i(t)$ — over Θ_i in multiples of standard deviation σ_{c_i} . Corresponds to a "reverse" cumulative distribution function of $c_i(t)$. Black solid lines indicate rCDF per kernel $k_i(t)$. Red dashed line indicates rCDF for a normal distribution with $\mu = 0$ and $\sigma = 1$.




Fig. 19: Relation between threshold value and pure-noise feature value for Fig. 8. Proportion of pure-noise kernel response $c_i(t)$ that exceeds threshold value Θ_i — which determines the average value μ_{f_i} of feature $f_i(t)$ — over Θ_i in multiples of standard deviation σ_{c_i} . Corresponds to a "reverse" cumulative distribution function of $c_i(t)$. Black solid lines indicate rCDF per kernel $k_i(t)$. Red dashed line indicates rCDF for a normal distribution with $\mu = 0$ and $\sigma = 1$.




Fig. 20: Relation between threshold value and pure-noise feature value for Fig. 9. Proportion of pure-noise kernel response $c_i(t)$ that exceeds threshold value Θ_i — which determines the average value μ_{f_i} of feature $f_i(t)$ — over Θ_i in multiples of standard deviation σ_{c_i} . Corresponds to a "reverse" cumulative distribution function of $c_i(t)$. Black solid lines indicate rCDF per kernel $k_i(t)$. Red dashed line indicates rCDF for a normal distribution with $\mu = 0$ and $\sigma = 1$.




Fig. 21: Relation between threshold value and pure-noise feature value for Fig. 10. Proportion of pure-noise kernel response $c_i(t)$ that exceeds threshold value Θ_i — which determines the average value μ_{f_i} of feature $f_i(t)$ — over Θ_i in multiples of standard deviation σ_{c_i} . Corresponds to a "reverse" cumulative distribution function of $c_i(t)$. Black solid lines indicate rCDF per kernel $k_i(t)$. Red dashed line indicates rCDF for a normal distribution with $\mu = 0$ and $\sigma = 1$.




Fig. 22: Threshold-dependent intensity invariance of species-specific feature sets. Same processing as in Fig. 8, using different kernel-specific threshold values Θ_i (multiples of pure-noise standard deviation σ_{η_i} of $c_i(t)$ for $\alpha = 0$). See also appendix Fig. 19). Average value μ_{f_i} of each feature $f_i(t)$ over α .