Distributions, then sound

Eleutherodactylus Research Studio · Fall 2026

Tom Devitt

2026-09-09

Species voices

You each spent the long weekend hunting for a recording of your claimed frog. Hands up:

- Found a published or archived recording
- Found nothing, and wrote that down

For a good share of our 40-odd species the honest answer is “nothing.” **If you searched, wrote down where you looked, and found no recording, you did the assignment.** A species with no public recording has no call description anyone can check, and that gap is the opening for your team’s paper. By the end of today you’ll know how a call gets turned into numbers, and how those numbers get summarized once you have a pile of them.

Two halves

1. Summarizing measurements

Last Wednesday every point was one frog. A team dataset will hold dozens to hundreds of calls from a handful of males, and the question changes from “where is this point” to “what is typical, and how much does it vary.” That is a *distribution*.

2. Sound

What a recording is, the three ways to picture it, and Köhler’s vocabulary for the pieces of a call. Monday you measure calls yourself; today you learn to look at one.

The same practice table

```
set.seed(206)
calls <- tibble(
  species = rep(c("E. marnockii", "E. cystignathoides", "E. guttilatus"),
                each = 40),
  temp_c  = round(runif(120, 17, 29), 1),
  call_rate_per_min = round(rep(c(8, 14, 5), each = 40) +
                            0.9 * (temp_c - 17) + rnorm(120, 0, 1.8), 1),
  dominant_khz = round(rep(c(2.9, 4.1, 3.4), each = 40) +
                       0.04 * (temp_c - 17) + rnorm(120, 0, 0.12), 2)
)
```

Still simulated, still one row per recorded male. Forty per species is a luxury you will not have: real call descriptions in this genus rest on 2 to 10 males, and that changes how you should draw them. Keep that in mind on every slide that follows.

A histogram

```
ggplot(calls, aes(x = dominant_khz)) +
  geom_histogram(binwidth = 0.1)
```

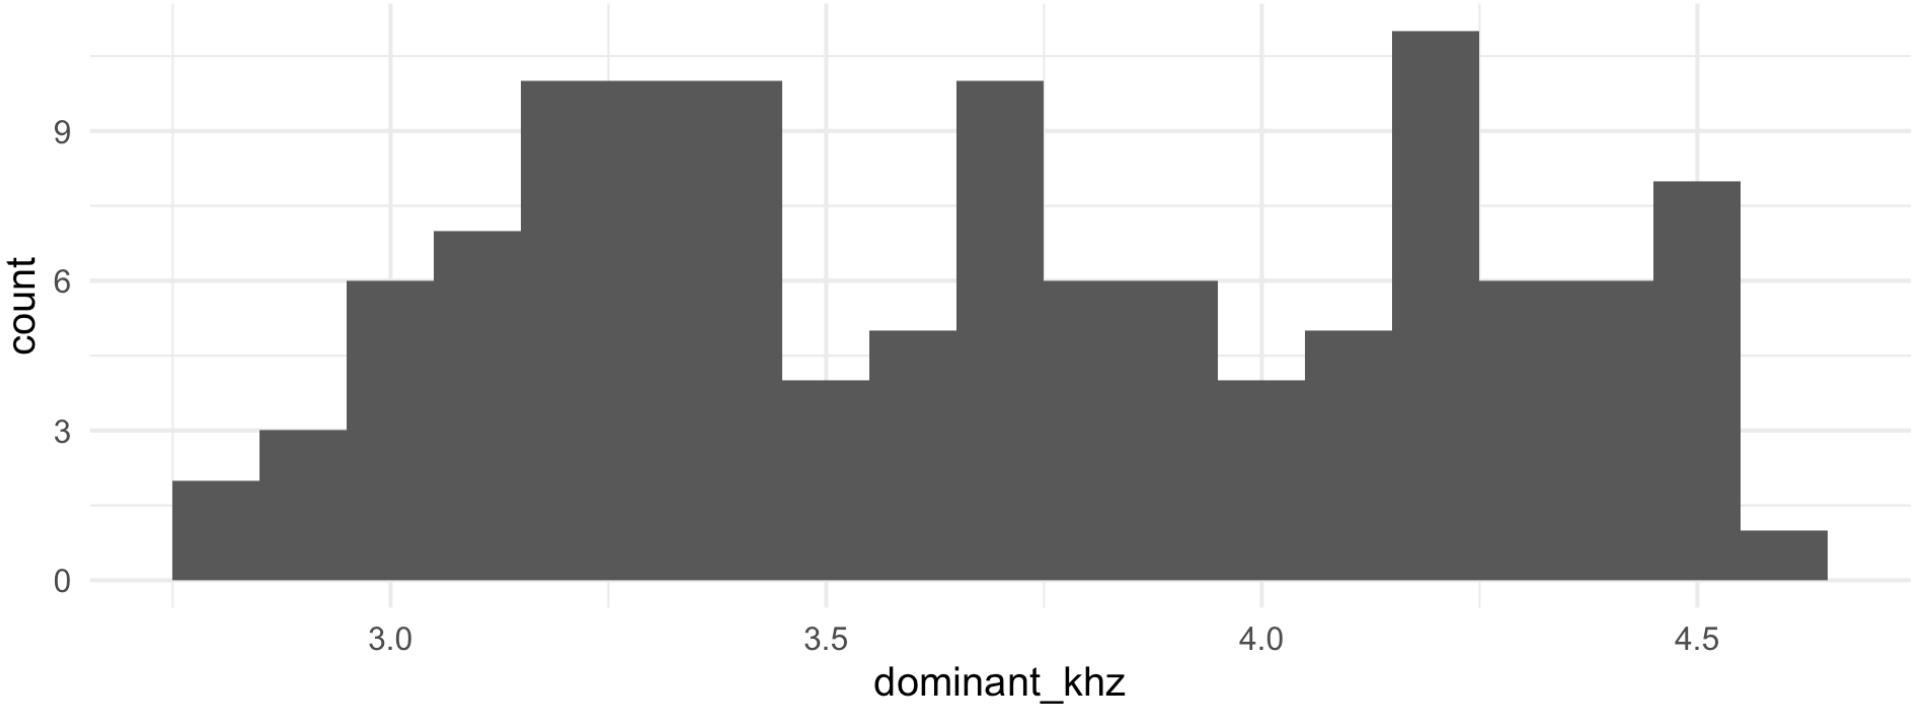

Only an x this time. ggplot2 chops the axis into bins and counts the rows in each. The height is a count, so y is computed for you. Three humps: three species.

Bin width is a choice you make

▶ Code

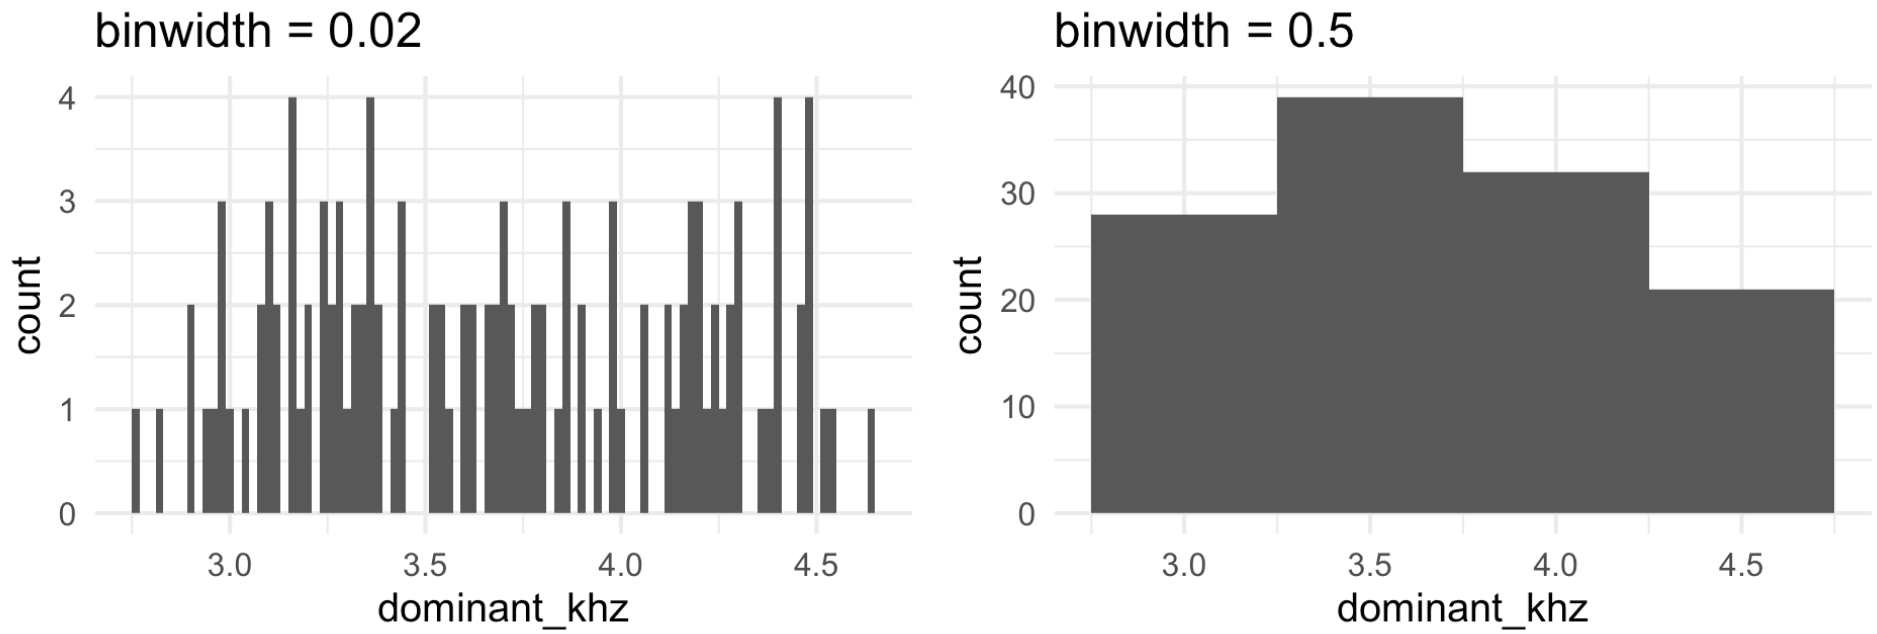

Too narrow and every bar is noise; too wide and the three species merge into one. ggplot2 warns you if you leave `binwidth` out, and you should take the hint. Try a few widths and keep the one that shows the structure without inventing any.

Groups: facets, not overlap

```
ggplot(calls, aes(x = dominant_khz)) +
  geom_histogram(binwidth = 0.1) +
  facet_wrap(~ species, ncol = 1) +
  theme(strip.text = element_text(face = "italic"))
```

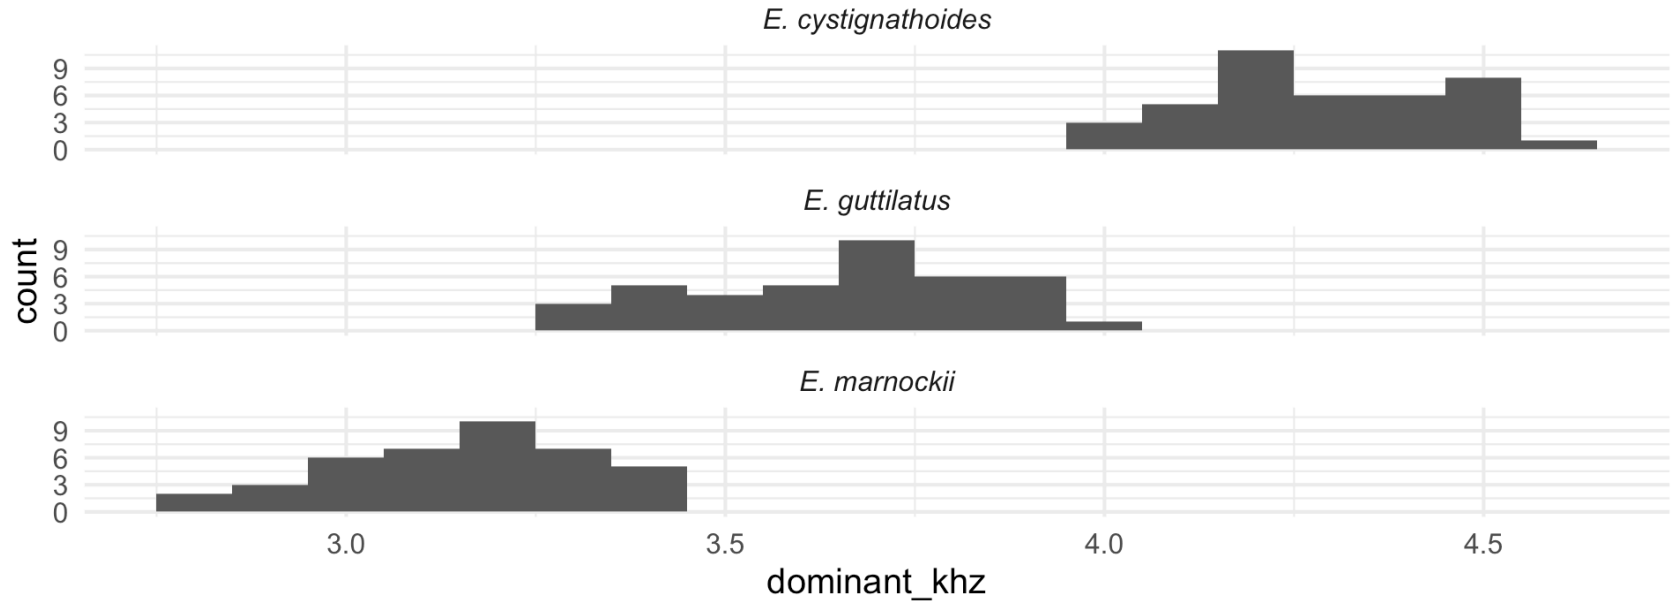

Stacking histograms in colors on one panel is hard to read once the groups overlap. Facets keep one axis and put the species under each other, so your eye compares positions.

Boxplots: five numbers per group

```
ggplot(calls, aes(x = species, y = dominant_khz)) +
  geom_boxplot() +
  theme(axis.text.x = element_text(face = "italic"))
```

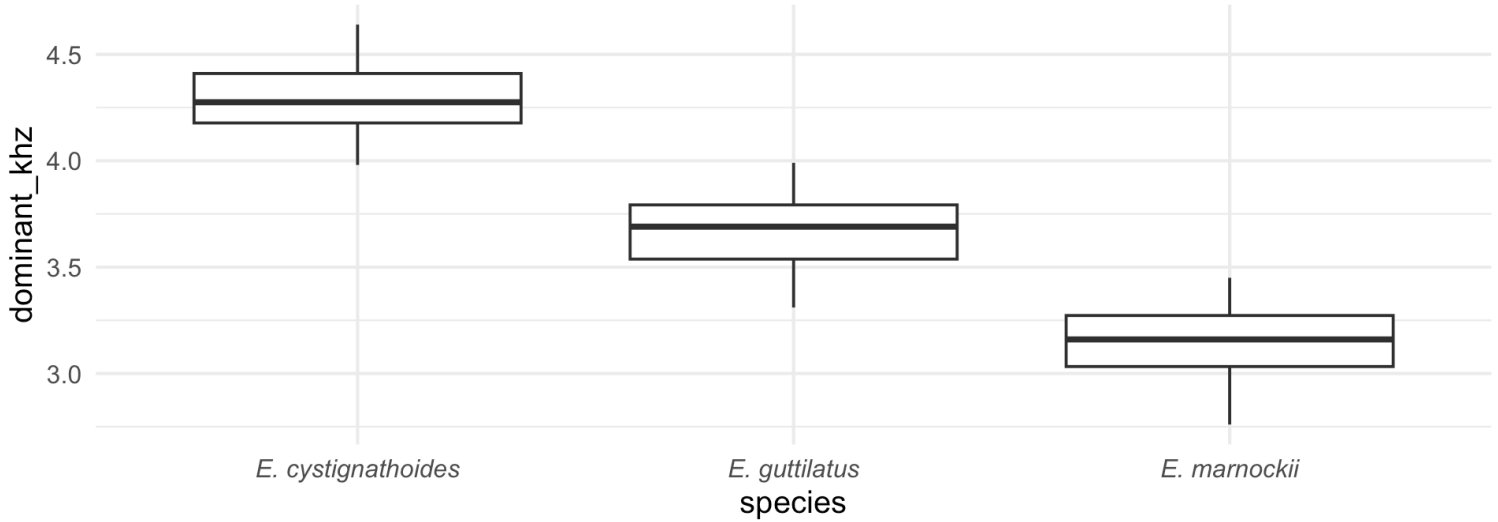

The line is the median. The box holds the middle half of the data (25th to 75th percentile). The whiskers reach the most extreme points within 1.5 box lengths, and anything past that is drawn as its own dot. A boxplot is a compressed histogram, and it needs enough points to compress, at least 5 per group (Krzywinski & Altman 2014, *Nature Methods*). Below that the box summarizes almost nothing and can mislead.

Show the frogs

```
ggplot(calls, aes(x = species, y = dominant_khz)) +
  geom_boxplot(outlier.shape = NA) +
  geom_jitter(width = 0.15, alpha = 0.5) +
  theme(axis.text.x = element_text(face = "italic"))
```

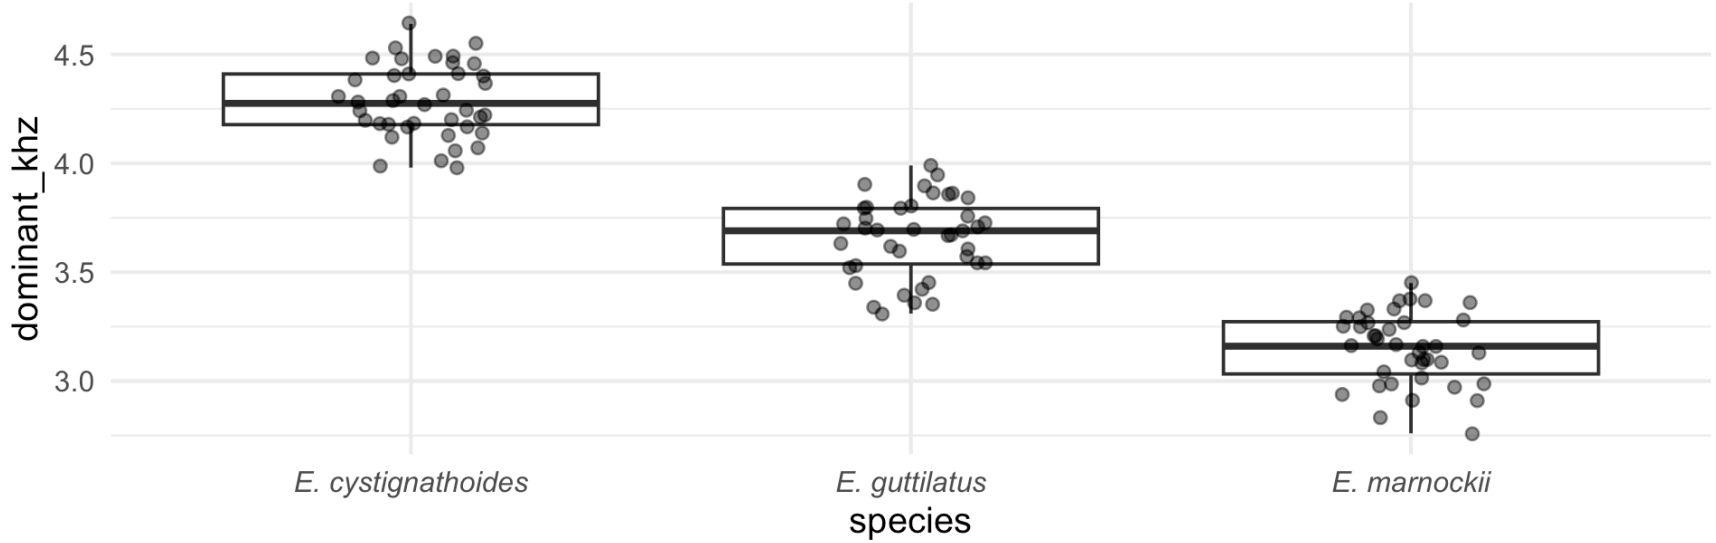

With 2 to 10 frogs per species, the points *are* the data, so draw them. `geom_jitter()` spreads the dots sideways so they don't stack. Keep the box for any group with 5 or more points; with fewer, drop it and show the points alone or with the mean. (`outlier.shape = NA` stops the boxplot drawing those frogs a second time.)

The numbers behind the picture

```r
calls |>
  summarise(
    n = n(),
    mean = mean(dominant_khz),
    sd = sd(dominant_khz),
    min = min(dominant_khz),
    max = max(dominant_khz),
    .by = species
  )
```

```
# A tibble: 3 × 6
  species                n  mean    sd   min   max
  <chr>              <int> <dbl> <dbl> <dbl> <dbl>
1 E. marnockii          40  3.15 0.164  2.76  3.45
2 E. cystignathoides    40  4.29 0.165  3.98  4.64
3 E. guttilatus         40  3.66 0.181  3.31  3.99
```

This is the table in every call description: mean ± SD (range), *n*. Köhler et al. ask for exactly these, per parameter, per species. With *n* of 2 or 3 the SD says little, so report the range and the individual values too. The plot is how you check the table isn't hiding something; the table is how you report it.

Half two: what is sound

A calling frog pushes air. The pressure at your ear, or at a microphone, rises and falls, and the microphone writes down that pressure 44,100 times a second (the *sample rate*). A recording is a very long column of numbers.

Two properties of the wave carry most of what we care about:

- **Frequency**: how many cycles per second, in hertz. We hear it as pitch. Frog calls live between roughly 0.5 and 8 kHz
- **Amplitude**: how big the swings are. We hear it as loudness

Everything on the next slides is a way of drawing that column of numbers.

Three views of one call

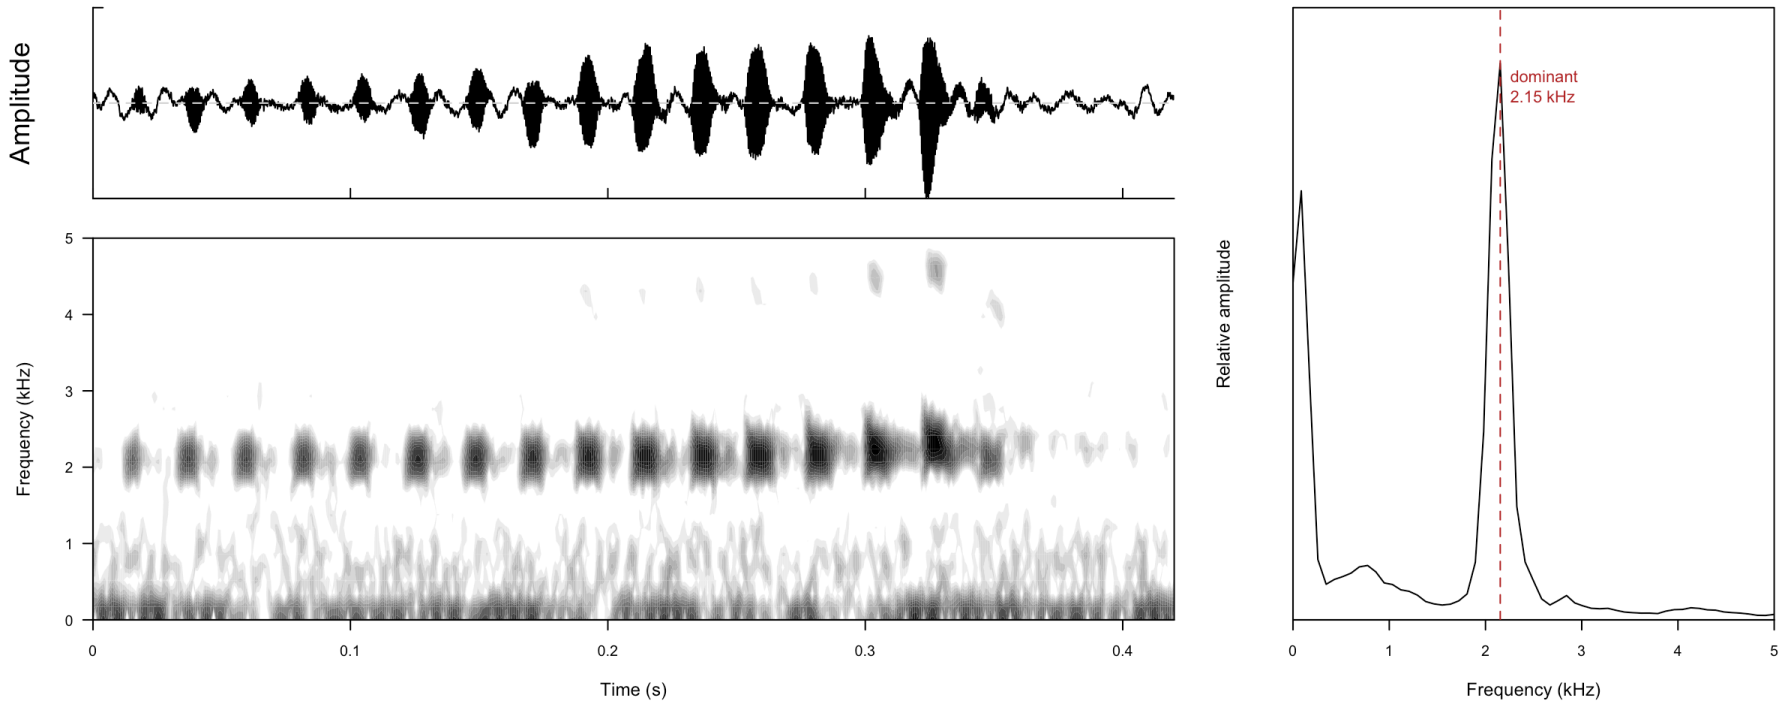

E. marnockii, Brackenridge Field Lab, one call, 0.42 s. Top left, the **oscillogram**: amplitude against time, the raw column of numbers. Bottom left, the **spectrogram**: time across, frequency up, darkness for energy. Right, the **power spectrum**: energy against frequency for the whole call, with the loudest frequency marked.

Reading a spectrogram

The spectrogram is a stack of short power spectra, one per slice of time. How long a slice is changes the picture:

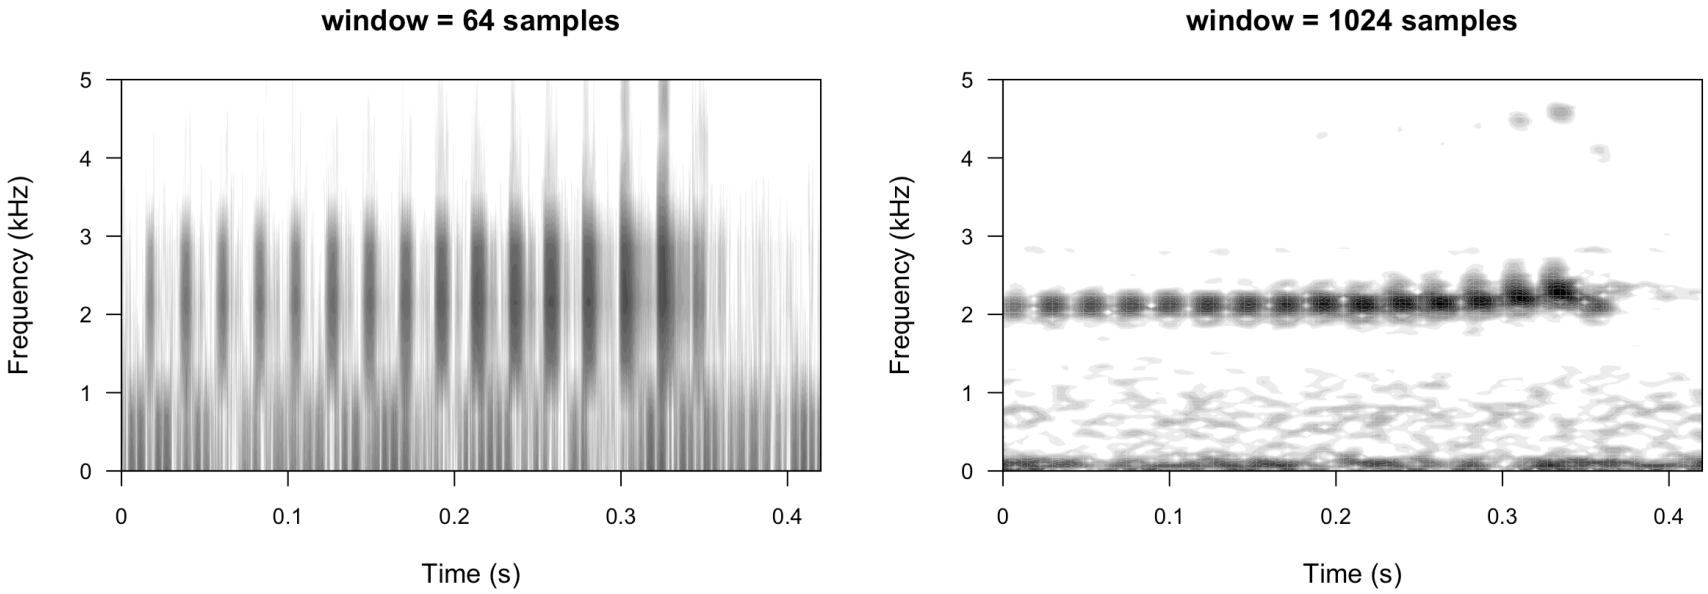

Short windows resolve time (each pulse is a crisp bar) and smear frequency. Long windows resolve frequency (a tight band at 2.1 kHz) and smear time. There is no free setting. Raven and warbleR both let you choose; report what you chose.

Köhler's vocabulary

Köhler et al. (2017) ask us to name the pieces of a call the same way, so that descriptions can be compared. From the bottom up:

- **Pulse**: a single burst of energy within a note, separated from its neighbors by a drop in amplitude
- **Note**: a continuous unit of sound, itself either tonal or made of pulses
- **Call**: the whole utterance a frog produces between silences, made of one or more notes

The catch is deciding where a note ends and the next call begins, which is often a judgment about the silences. Köhler calls the two defensible options *note-centered* and *call-centered* schemes, and asks you to say which one you used. We'll use call-centered unless a species forces otherwise.

A call series

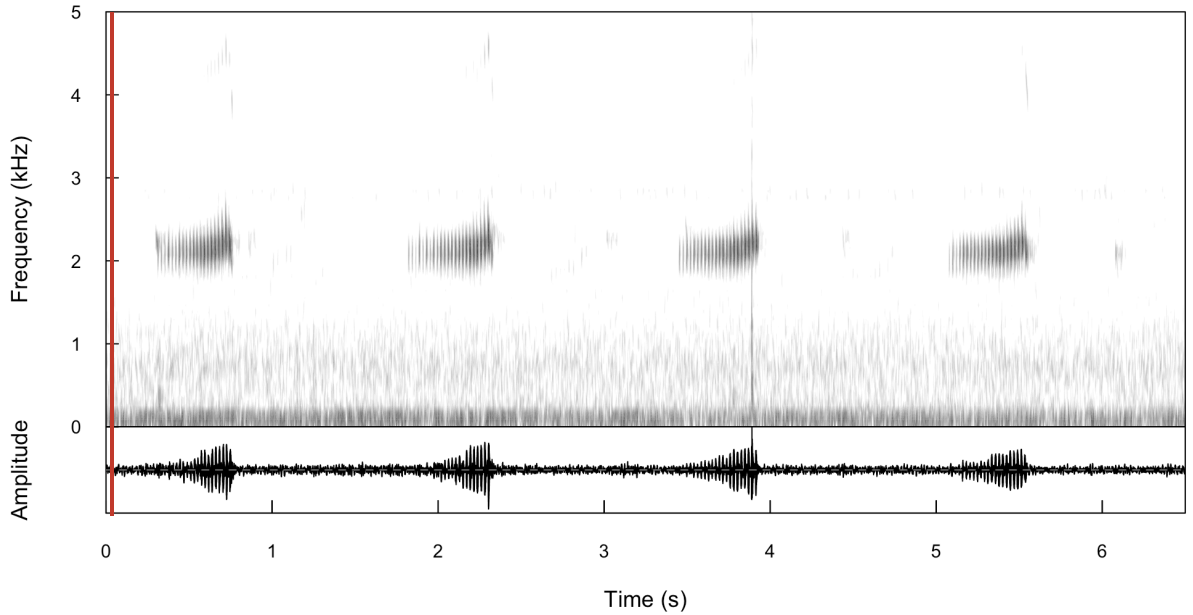

0:00 / 0:06 play at ¼ speed

Six and a half seconds of the same male. Four calls, each a short pulsed note with rising amplitude, a little above 2 kHz; between them, insect noise below 1 kHz and one loud click. Press play and the red line follows the sound. Click the figure to jump. At quarter speed the pitch drops two octaves and the pulses become countable by ear.

Read the numbers off it

From the two figures, by eye:

parameter	this frog	how you got it
call duration	about 0.3 s	oscillogram, first to last pulse
inter-call interval	about 1.6 s	spectrogram, start to start
call rate	about 37 per minute	60 / interval
pulses per call	about 15	count the bars
pulse rate	about 45 per second	pulses / duration
dominant frequency	2.15 kHz	power-spectrum peak

Your teams will measure these for every call in every recording. Each becomes a column, and each column gets a histogram and a mean ± SD (range) with *n*.

Which numbers you can trust

- **Temperature matters.** Call rate and pulse rate climb with air temperature in anurans generally, dominant frequency a little. A measurement without the temperature at the time is hard to compare with anything. Record it; recover it from field notes when you can
- **Compression matters.** The AudioMoth recordings are WAV. Many of the older handheld recordings you'll use are MP3, which throws away the quiet parts and the fine frequency detail. From MP3, stick to durations, rates, counts, and dominant frequency. Skip amplitude and fine spectral shape
- **Which frog matters.** One male recorded twenty times is $n = 1$ male, not twenty. Köhler asks for both counts

Your turn: worksheet

On the EDU pod:

```
file.copy("/stor/work/FRI_Bio206_TD_Spring2026/05-distributions-worksheet.qmd", "~/")
```

Five tasks on today's `calls` table (histograms, bin widths, facets, a boxplot with points, a summary table), then a sixth where you read the *marnockii* figures the way we just did. Render it and **submit the HTML to Canvas**.

Monday

- **Bioacoustics II, hands-on.** We open a recording in Raven Lite and in warbleR on the pod and measure a call together. Bring headphones
- **Read before Monday:** Köhler et al. 2017, the *Methods* section on spectrogram settings and the *Terminology* section in full this time (link on Resources). Skim the parameter list at the end; you'll be picking from it
- **Wed Sep 16** is a work session: teams start measuring their own species. Team pages and materials are already on the pod share

Alt text added with Claude Code (Claude Opus 5.5, Anthropic).