Fourier Analysis Worksheet + Hit Song Statistics Problem Set

Part 1: Fourier Analysis in Practice — A Hands-On Guide

This section walks you through real Fourier analysis of audio files using two free, powerful tools: **Sonic Visualiser** (advanced music analysis) and **Audacity** (simple spectrum plotting).

🧰 Software & Setup

Tool	Download	Best for
Sonic Visualiser	sonicvisualiser.org	Spectrograms, pitch tracking, Vamp plugins, detailed music analysis
Audacity	audacityteam.org	Quick spectrum plots (Analyze → Plot Spectrum), EQ analysis, waveform editing
Vamp Plugins (optional)	vamp-plugins.org	Automated feature extraction: tempo, chroma, note detection

> 💾 **Example audio files:** Use any short MP3 or WAV file (e.g., a pop song chorus, a pure tone you record, or the sample files provided with Sonic Visualiser). For best results, use files with clear, isolated sounds (single instrument or vocal).

🎯 Step-by-Step Exercise 1: Viewing a Spectrogram in Sonic Visualiser

Goal: Understand how a song's frequency content changes over time.

1. **Launch Sonic Visualiser** → *File* → *Open* → select your audio file (MP3/WAV).
2. **Add a spectrogram pane:** *Pane* → *Add Spectrogram* (or click the spectrogram icon).
3. **Adjust display:** Right-click the spectrogram → *Colour Scale* → choose *Linear* or *Log*.
4. **Explore:**
   - A **spectrogram** is a time-frequency breakdown — the output of a series of short-time Fourier transforms, one for each horizontal step.
   - Dark red/yellow = strong energy at that frequency.
   - Hover over the spectrogram to see exact frequency and time.
5. **Compare:** Add a *Waveform* pane next to the spectrogram. See how dense rhythmic sections appear as vertical streaks.
6. **Save your analysis:** *File* → *Export Image File* to capture a screenshot for your notes.

📝 **Journal Question:** Pick a 5-second clip. What frequencies dominate? Do you see harmonics (vertical stripes at multiples of the fundamental)? Which instrument do you think is responsible?

🎯 Step-by-Step Exercise 2: Plotting a Frequency Spectrum in Audacity

Goal: Quantify which frequencies are present in a selected sound.

1. **Launch Audacity** → *File* → *Import* → *Audio* → select your file.
2. **Select a short, steady portion** of the audio (e.g., a held vocal note or a drum hit).
3. **Plot the spectrum:** *Analyze* → *Plot Spectrum*.

> **What's happening?** Plot Spectrum takes the selected audio (a set of sound pressure values at points in time) and converts it to a graph of frequencies against amplitudes using a Fast Fourier Transform (FFT). All the values are averaged to create the graph.

4. **Adjust parameters for better resolution:**
   - *Size:* Larger = better frequency resolution, but requires more audio samples. For a 44,100 Hz file, a size of 4096 gives ~10.8 Hz resolution.
   - *Function:* Leave as *Hann* (default). This controls how the data is windowed.

- *Axis:* Linear frequency vs. Log frequency — choose Log for musical perception.

5. **Read the plot:** Move your cursor over peaks — the status bar shows exact frequency (Hz) and amplitude (dB). The highest peak is usually the fundamental pitch.
6. **Export the graph:** In the Plot Spectrum window, click *Export* to save data as a text file for further analysis.

📝 **Journal Question:** What is the strongest frequency in your selection? Does it match the expected note? Are there clear harmonics at ~2×, 3×, 4× that frequency?

🎯 Step-by-Step Exercise 3: Comparing Two Different Chords / Sections

Goal: Quantify harmonic simplicity and energy distribution.

1. In **Audacity**, select two different 1-second sections (e.g., a quiet intro vs. a loud chorus).
2. Run *Plot Spectrum* for each.
3. **Compare:**
    - Which section has more high-frequency content?
    - Which spectrum looks cleaner (fewer peaks)?
    - Use the *Export* button to save both spectra as .txt files and compare the numbers.

📝 **Journal Question:** Does the louder section also have stronger harmonics relative to the fundamental? What does that tell you about the arrangement?

🎯 Step-by-Step Exercise 4: Automated Feature Extraction with Vamp Plugins (Optional)

Goal: Let the computer count rhythmic periodicity and tonal features.

1. In **Sonic Visualiser**, install *Vamp plugins* (free from vamp-plugins.org). Recommended: *Queen Mary* plugin set (tempo, beats, chroma).
2. Go to *Transform → Available Transforms*.
3. Choose **"Tempo and Beat Tracker"** → hit *Run*.
4. A new layer will show detected beats and estimated BPM.
5. Compare the computer's BPM to the actual song. Is it accurate?
6. Try **"Chromagram"** (pitch class profile) — this shows how notes cycle through the 12 pitch classes over time.

> **Why this matters:** These plugins calculate high-level features directly from the short-time Fourier transform. Exactly the same techniques used in hit prediction research.

📝 **Journal Question:** How close is the detected BPM to a round number? Is the tempo stable or does it fluctuate?

🎯 Step-by-Step Exercise 5: Testing Phase Cancellation (Bonus)

Goal: Use Fourier analysis to detect interference.

1. Create two identical sine waves in separate tracks in Audacity (Generate → Tone).
2. Invert one (*Effect → Invert*).
3. Mix them (*Tracks → Mix → Mix and Render*).
4. Select a section and run *Plot Spectrum*.
5. Result: Almost nothing — total cancellation at all frequencies! This shows destructive interference in the frequency domain.

📝 **Journal Question:** Why does the Plot Spectrum show near-zero amplitude even though the original tones were present? Connect this to the *Calculus of Harmony* video.

Part 2: Hit Song Statistics — Problems Using Real Data

These problems use real statistics drawn from academic studies of Spotify chart data. No coding required — just a calculator and your reasoning.

> **Reference dataset:** A 2024 study collected 5,010 distinct hit songs and 6,699 viral songs from Spotify’s charts, analyzing intrinsic features (acoustic, lyrical) and extrinsic features (artist, chart history).

🎵 Problem Set 1: Describing Hit Song Features

Use the following summary statistics from a 2024 analysis of Spotify top tracks:

Feature	Mean	Median	Range
BPM (tempo)	122.5	121	65 – 206
Energy (%)	64.3	66	9 – 97
Danceability (%)	67.0	69	23 – 96
Valence (positivity, %)	51.4	51	4 – 97

Q1: Based on the mean and median, are these distributions roughly symmetric or skewed? Hint: Compare mean vs. median for each feature.

Q2: A new song has BPM 135, Energy 75, Danceability 85, Valence 60. How does it compare to the average hit song? Is it above average in all categories?

Q3: Which feature shows the widest range? What might this tell you about hit song variety?

Q4: Valence measures positivity (high = happy, low = sad). The mean is ~51%. What does this imply about the emotional balance of hit songs?

🎵 Problem Set 2: Comparing Hit vs. Viral Songs

Researchers compared *hit songs* (charting) vs. *viral songs* (trending on social media). They found that intrinsic audio features (acousticness, danceability, energy, etc.) can distinguish between them.

Q5: A researcher claims: "Danceability is a stronger predictor of viral success than energy." If you had average danceability scores of 68% for hits and 72% for viral songs, which group is more danceable? What does this suggest?

Q6: The loudness of hit songs increased by approximately 2.5 dB per decade from 1980 to 2020. If a 1980 hit measured 85 dB, what would a 2020 hit measure? How many times more intense is the 2020 hit? (Hint: $dB = 10 \log_{10}(I/I_0)$)

🎵 Problem Set 3: Predictive Models

A 2024 prediction study using Spotify's U.S. Top 200 Daily Charts built a dataset of 14,639 unique songs and achieved **97% accuracy** predicting chart success using Random Forest and XGBoost models.

Q7: If the model predicts 100 new songs, how many would it correctly classify as hit or not (assuming the 97% accuracy holds)?

Q8: Even when stream count and rank history were excluded, models trained solely on audio attributes retained predictive power. Why is this significant for independent artists?

Q9: Another study found that Random Forest improved prediction accuracy by 7.1% compared to average scores. If the average model scored 78% accuracy, what would the Random Forest accuracy be?

🎵 Problem Set 4: Streaming Distribution

Streaming success is highly concentrated: most tracks accumulate fewer than 1 billion streams, while only a few surpass this number.

Q10: Interpret this histogram description: "The sharp decline after the first few bins indicates that streaming success is highly concentrated among a small number of songs." What does this imply for a new artist trying to chart?

Q11: If the top 1% of songs receive 50% of total streams, and there are 10,000 songs on the platform, how many songs account for half of all streams? How many streams would each of those top songs have on average (assuming total streams = 100 billion)?

Bonus: A study examining the first 200 plays of a new song could predict its chances of becoming a hit (top 1% most played). If a song has 200 plays and is in the top 1%, what percentile of popularity does that represent?

🎵 Problem Set 5: Acoustic Feature Analysis

Spotify's API provides numerical audio features for every track. Table 1 in Oliveira et al. (2024) lists:

Feature	Description
acousticness	Probability (0–1) of being acoustic
danceability	Probability (0–1) suitable for dancing
energy	Intensity (0–1) — perceived loudness, timbre, entropy
instrumentalness	Probability (0–1) — no vocals
liveness	Probability (0–1) — audience present
speechiness	Probability (0–1) — spoken words
valence	Positiveness (0–1) — high = happy
tempo	Beats per minute (BPM)
loudness	Measured in decibels (dB)

> Source: Table 1 in.

Q12: A song has energy = 0.85, danceability = 0.75, valence = 0.90. Categorize its mood and likely genre.

Q13: If a hit song has acousticness = 0.95 and instrumentalness = 0.02, is it likely an acoustic ballad or a spoken-word piece? Explain.

Q14: Two songs: Song X has loudness = –6 dB, Song Y has loudness = –3 dB. How many times more intense is Song Y? (Hint: ΔdB = 20 $\log_{10}(p_2/p_1)$ for pressure; for power, use 10 $\log_{10}$.)

Q15: Using the formula **centroid = Σ(frequency × amplitude) / Σ(amplitude)**, estimate the spectral centroid of a sound with the following partials: 440 Hz at 0 dB, 880 Hz at –6 dB, 1320 Hz at –12 dB. Would this sound be described as "bright" or "dark"?

📎 Answer Key

Part 2 Answers

Q#	Answer
1	Energy: mean 64.3 vs median 66 → slight left skew (more low-energy songs). Danceability: mean 67.0 vs median 69 → also left skew. Valence: mean 51.4 vs median 51 → nearly symmetric.
2	Above average in all: BPM (+12.5 above mean), Energy (+10.7), Danceability (+18.0), Valence (+8.6).
3	Valence range (93) is widest. Hit songs vary greatly from very sad to very happy. Energy range (88) also wide.
4	Hits are emotionally balanced — roughly equal numbers of positive and negative songs.
5	Viral songs are more danceable (72% > 68%). Danceability may be more important for social media engagement than pure energy.
6	2.5 dB/decade × 4 decades = +10 dB. 2020 hit = 95 dB. Intensity ratio = 10^(10/10) = 10 times more intense.
7	97 out of 100 correctly classified (assuming no overfitting).

Q#	Answer
8	Independent artists can make data-driven decisions about song structure and production without needing existing popularity.
9	Random Forest accuracy = 78% + (7.1% of 78) = 78% + 5.54% = **83.54%**.
10	Extremely difficult to reach top tier; most songs get relatively few streams. Implies importance of playlist placement and algorithmic discovery.
11	Top 1% of 10,000 = 100 songs account for 50% of streams. Each of these 100 averages (0.50 × 100B)/100 = 500M streams.
12	High energy + high danceability + high valence → likely upbeat pop or dance-pop track. Happy, energetic, and rhythmic.
13	High acousticness (95% likely acoustic) + low instrumentalness (98% likely has vocals) → acoustic ballad with singing, not spoken word.
14	ΔdB = −6 − (−3) = −3 dB. Intensity ratio = 10^(−3/10) = 10^(−0.3) ≈ 0.5. Song Y is twice as intense as Song X.
15	Centroid = (440×1 + 880×0.5 + 1320×0.25)/(1 + 0.5 + 0.25) = (440 + 440 + 330)/(1.75) = 1210/1.75 ≈ 691 Hz. Relatively low → "dark" sound (most energy in low mids).

📌 Glossary of Key Terms

Term	Definition
Fourier Transform	Mathematical conversion of a time-domain signal into its frequency components.
FFT (Fast Fourier Transform)	Efficient algorithm for computing the Fourier transform on digital audio.
Spectrogram	Visual representation of frequency content over time.
Spectral Centroid	Weighted average of frequencies; high = bright timbre, low = dark timbre.
Chroma Features	12-dimensional representation of pitch class distribution.
Vamp Plugins	Extensible plugins for automated feature extraction in Sonic Visualiser.
Intrinsic Features	Song characteristics derived from the audio itself (not external metadata).

End of Worksheet

Created for use with "Statistical Analysis of Pop Hits" explainer video.