



ACHOS

FOR THE GUARDIANS OF SOUND

ἄχος · Homeric Greek · The ache of one who refuses compromised sound

TECHNICAL MANUAL

MANUAL	APP	PLATFORM	DEVELOPER	CODING AGENT	DATE
v1.0	v1.0.0	macOS · Apple Silicon	Martin Gallardo	Claude AI · Anthropic	June 2026



ACHOS

FOR THE GUARDIANS OF SOUND

ἄχος · Homeric Greek · The ache of one who refuses compromised sound

TECHNICAL MANUAL

MANUAL	APP	PLATFORM	DEVELOPER	CODING AGENT	DATE
v1.0	v1.0.0	macOS · Apple Silicon	Martin Gallardo	Claude AI · Anthropic	June 2026

Version History

Version	Date	Milestone	Note
v0.1.0	Apr 24 2026	First playback	Custom DSD decoder. Achos brand established. First music file played.
v0.2.0	Apr 30 2026	Audio path	Hog Mode, BitPerfect mode, DoP (DSD over PCM) for DSD64/DSD128.
v0.3.x	May 12 2026	Stability	Audio path hardening. Race conditions and DAC timing fixes.
v0.4.0	May 18 2026	Exclusive path	AVAudioEngine exclusive path. Rate separation for shared vs. exclusive.
v0.5.x	May 19 2026	Spectral engine	Dual-FFT architecture. Fingerprint export. Full Analysis engine.
v0.6.0	May 26 2026	Discogs	Discogs OAuth 1.0a integration. Collection sync. Pressing Intelligence.
v0.7.x	May 28 2026	Stability II	Keychain hardening. Security-scoped bookmarks. EBU R128 loudness gating.

★ FROM v0.8.0 — ACHOS BECOMES AN AI-AUGMENTED INSTRUMENT

Version	Date	Milestone	Note
v0.8.0	May 30 2026	★ AI Filemaster	Gemini 2.5 Flash forensic engine. Spectral verdict system. First AI integration.
v0.8.1	May 31 2026	True Peak fix	Critical sinc FIR kernel correction. Gain error resolved.
v0.9.0	Jun 2 2026	★ GRaiL	AI musical narrative. Gemini with mandatory web search. Guardian's Path gating.
v0.9.1	Jun 5 2026	Stabilization	18 bug fixes. Vinyl Oracle. Guardian's Path dot indicator system.
v0.9.2	Jun 5 2026	Guardian's Path	Unified window. Gold AI highlight system. Genre-aware Filemaster. Volume monitor.
v1.0.0	Jun 8 2026	Public release	Onboarding. OAuth fix. Verdict taxonomy. Menubar polish. App Store ready.

★ Marks AI-augmented features. All versions developed by the Human-AI Co-Creation team: Martin Gallardo (AI-Augmented Developer) and Claude AI · Anthropic (Coding Agent).

Part I — Introduction

1.1 What is Achos

Achos is a macOS music player built exclusively for Apple Silicon, designed for those who treat their music library as a serious collection rather than a background service. The name derives from the Homeric Greek word ἄχος (áchos) — the deep ache felt by one who cares deeply and refuses to accept less than what is right. For the audiophile, that ache is the refusal to accept compromised sound.

Achos was conceived from a single, uncompromising statement at v0.1.0: **listen to local music.** Everything built since then — the spectral analysis engine, the AI forensic system, the Discogs integration, the GRaiL (Guided Research for Audiophile Insight and Listening) narrative feature — serves that original mission. The listener comes first. The technology serves the music.

1.2 Core philosophy

Listening first. The playback engine is sacred. No analysis, scanning, or AI task ever runs on the audio render thread. Music plays without interruption, always.

Honest measurement. Achos applies internationally recognised standards — ITU-R BS.1770, EBU R128 — to every file it analyses. The results are not estimates. They are measurements, reported with the same precision used in professional broadcast facilities.

Informed listening. A file's format claim is not a guarantee. Achos tells the listener what their files actually contain — not what the filename or tag says.

1.3 Who this manual is for

This manual is written for anyone who loves music and wants to understand it more deeply — whether an enthusiastic listener who has just discovered hi-res audio, or a seasoned audio engineer. Technical sections are clearly marked. Plain-language summaries always precede the technical detail.

A note on limitations. Achos is not a certified laboratory instrument. Its measurements and forensic verdicts are informative references, not legal or engineering certifications. Results should be weighed in context. Spectral inference is probabilistic, not deterministic — the same measurement can have more than one valid explanation.

1.4 How to read this manual

The manual is organised in six parts plus appendices. Readers new to hi-res audio are encouraged to read Part II before Part III. Part V contains the full technical methodology. Appendix C documents the COFI-based Human-AI Co-Creation process used to develop Achos.

Part II — Understanding Your Audio Files

2.1 What hi-res audio actually means

A digital audio file captures sound by taking thousands of measurements of the audio waveform every second. The **sample rate** determines how many measurements are taken per second (expressed in kHz), and the **bit depth** determines how precisely each measurement is recorded.

CD-quality audio uses 44.1 kHz / 16-bit — 44,100 measurements per second, each with 65,536 possible values. Hi-res audio typically means 96 kHz or 192 kHz at 24-bit. In practice, what matters is whether the file actually contains genuine high-resolution information, or whether it is simply a lower-resolution recording repackaged in a larger container.

DSD (Direct Stream Digital) is a fundamentally different approach. Instead of measuring amplitude thousands of times per second, DSD encodes audio as a rapid stream of single bits at 2.8 MHz or higher — 64 times the CD sample rate. DSD64 files carry the .dsf or .dff extension and are the native format of the Super Audio CD (SACD).

2.2 Key concepts — plain language

Term	Plain-language definition
Decibels (dB)	A logarithmic unit for measuring level. 0 dBFS is the digital maximum; everything below is negative. -6 dB = half amplitude. dBTP measures the peak after digital-to-analogue reconstruction.
LUFS	Loudness Units relative to Full Scale. Measures perceived loudness over time. Audiophile classical recordings typically reach -25 LUFS or quieter, preserving full dynamic range.
FFT	Fast Fourier Transform. Converts an audio signal into a frequency map — showing how much energy exists at each frequency. Achos uses FFT to detect upsampling, conversion, and genuine hi-res content.
Noise floor	The level of background noise inherent to a recording or format. DSD pushes quantisation noise into the ultrasonic range using noise shaping — a distinctive spectral signature Achos uses to verify DSD authenticity.
Crest factor	The ratio of peak to RMS level in dB. A high crest factor (>20 dB) indicates wide dynamic range.

2.3 Why format claims cannot always be trusted

A file's container format and extension describe how audio data is stored — not where it came from. A 192 kHz FLAC file may contain a genuine 192 kHz recording, or a 44.1 kHz recording upsampled and re-saved. Both files look identical in a standard player. Only spectral analysis reveals the difference.

Example. Norah Jones — *One Flight Down* exists in two hi-res versions. The 96 kHz version shows audio content stopping at 22 kHz — the fingerprint of a CD-quality source repackaged at a higher rate. The 192 kHz version extends naturally to 78.4 kHz, consistent with a genuine high-resolution transfer from the original master. Both files claim to be hi-res. Only one delivers it.

Part III — Achos Features

3.1 Main window — signal stream

The main window is the command centre of Achos. The audio signal follows a strict path from file to output:




Fig. 1 — Main audio signal path. Output may be the macOS system mixer (shared mode) or a connected DAC in exclusive Hog Mode.

In **Shared Mode**, audio passes through the macOS system mixer, which may apply sample rate conversion. In **Hog Mode with BitPerfect**, Achos claims exclusive control of the DAC and sends audio at the file's native sample rate with no conversion. DSD playback via DoP (DSD over PCM) requires Hog Mode exclusively.

3.2 Quick Analysis

Quick Analysis is the fastest way to assess a track. A single button press in the Quality Classification section of the main window triggers an immediate analysis of the selected file. Results appear within seconds and include format identification, sample rate, bit depth, and a quality classification with confidence level. Quick Analysis does not perform deep spectral FFT analysis — it classifies based on file format metadata and basic signal properties. Use Full Analysis for spectral verification.

3.3 Full Analysis

Full Analysis performs a complete examination of the audio file using the dual-FFT spectral engine, True Peak measurement, EBU R128 loudness gating, and brickwall cutoff detection. The result is saved as an Achos Fingerprint — a detailed record of the file's true spectral characteristics that persists in the local database and can be exported as a plain text report. Full Analysis is the first required step in the Guardian's Path.

3.4 AI Filemaster

AI Filemaster submits the Full Analysis fingerprint to an AI engine (Google Gemini 2.5 Flash) for forensic evaluation. The AI examines the spectral evidence and returns a structured verdict with supporting reasoning, evidence summary, and recommended action.

API key required. AI Filemaster requires a personal Google Gemini API key, which the user must obtain independently from Google AI Studio at aistudio.google.com. Achos does not provide, manage, or store API keys on behalf of users. The key is stored securely in the macOS Keychain. Users must comply with Google's Gemini API terms of service when using this feature.

Achos does not transmit any personally identifiable information — only the anonymised spectral fingerprint data is sent. Results are cached locally and never automatically re-run.

The four verdicts

Verdict	Meaning	Recommended action
Verified	All spectral markers consistent with claimed format.	Keep
Inconclusive	Ambiguous evidence — cannot determine provenance.	Review
Likely Upsampled	Content does not reach claimed resolution.	Consider replacing with a verified source
Likely Converted	Format does not reflect original recording medium.	Keep if quality satisfactory / Keep if hi-res origin is verified

Table 1 — AI Filemaster verdict taxonomy. Verdicts reflect spectral evidence only and are informative, not definitive.

3.5 GRaiL

GRaiL — the Guided Research for Audiophile Insight and Listening — is the most human feature in Achos. Where AI Filemaster examines *what* a file is, GRaiL explores *what it means*. It uses Google Gemini with mandatory web search to research the recording history, the artists, the studio, the cultural moment — and returns a warm, readable narrative that places the music in its full context. The result can be exported as a Markdown file for reference or sharing.

GRaiL requires Guardian's Path completion. Full Analysis and AI Filemaster must both complete successfully before GRaiL becomes available for a track. This ensures the narrative is always paired with verified provenance data.

3.6 Guardian's Path

Guardian's Path is the unified analysis journey for any track. It enforces a strict sequence — Full Analysis → AI Filemaster → GRaiL — ensuring each step builds on the verified output of the previous one.

● Full Analysis	● AI Filemaster	● GRaiL
Spectral fingerprint True Peak · EBU R128 Cutoff detection	Forensic verdict Evidence summary Recommended action	Recording history Artist context Cultural narrative
Green dot	Orange dot	Gold dot

Fig. 2 — Guardian's Path sequence. Steps must complete in order. Dot colors match those displayed in the library and main window.

3.7 Vinyl Oracle

The Vinyl Oracle draws from the user's Discogs collection and selects an album or track at random — an invitation to listen to something already owned but perhaps forgotten. Two modes: Pick Album and Pick Track.

3.8 Pressing Intelligence

Pressing Intelligence connects Achos to the user's Discogs collection and identifies the specific pressing of a physical release: country of origin, year, label, catalogue number, and pressing notes. Requires a Discogs account connected via OAuth 1.0a.

3.9 Equalizer and playback controls

Achos includes a parametric equalizer accessible from the View menu (⌘E) or the main window toggle. Keyboard shortcuts: Space (play/pause), ⌘. (stop), ⌘→ (next), ⌘← (previous).

Part IV — Reading and Saving Reports

4.1 The Achos Fingerprint (.txt)

The Achos Fingerprint is the primary output of Full Analysis. Every field is described below.

Section	Field	Description
FILE	Path	Full path to the audio file on disk.
	Container	File format (FLAC, DSD, MP3, etc.).
	File size	Total file size in MB.
METADATA	Title / Artist / Album	Tags embedded in the file. May be absent.
	Duration	Total playback duration.
TECHNICAL	Sample Rate	Samples per second (kHz). For DSD: 2.8 / 5.6 / 11.2 MHz.
	Bit Depth	Bits per sample. Always 1 for DSD files.
	Bit Rate	Data throughput in kbps.
QUALITY	Structure Type	Format category (Lossless, Lossy, 1-Bit Stream, etc.).
	Upsampled	Yes / No / Unknown — based on spectral evidence.
	Confidence	Low / Medium / High — reliability of classification.
LOUDNESS	Integrated Loudness	EBU R128 gated loudness in LUFS.
	True Peak	Maximum inter-sample peak in dBTP (ITU-R BS.1770).
	Noise Floor	Measured background noise level in dBFS.
	Crest Factor	Peak-to-RMS ratio in dB. Higher = more dynamic.
SPECTRAL	FFT size	Analysis window size.
	Nyquist	Half the sample rate — maximum theoretically present frequency.
	Detected cutoff	Frequency where brickwall filter was detected, or Not detected.
	HF extension	Highest frequency with measurable energy above noise floor.
	Subsonic energy	Yes / No — energy detected below 20 Hz.
OCTAVE BAND	31 Hz - 16 kHz	ISO 266 octave band energy levels in dBFS.

Table 2 — Achos Fingerprint field reference. Export via the fingerprint icon in the analysis window.

4.2 AI Filemaster forensic report (.md)

The Filemaster report contains: verdict and confidence level; one-sentence executive summary; evidence section listing key spectral measurements; detailed AI reasoning; conclusion; and forensic evidence summary. Suitable for sharing with other collectors or archiving alongside the audio file.

4.3 GRaiL narrative export (.md)

The GRaiL narrative exports as a Markdown file containing the full AI-generated text about the recording — its history, the artists, the production context, and its place in musical culture. Readable in any text editor, note-taking application, or Markdown renderer.

Part V — Measurement Methodology

5.1 Dual-FFT architecture

Achos employs two simultaneous FFT analysis passes on every file, each optimised for a different frequency range:

High-frequency pass: An FFT analysis at the file's native sample rate, with a Hann window and 50% frame overlap. Examines content up to the Nyquist frequency. Used for brickwall cutoff detection, DSD noise shaping curve analysis, and HF extension measurement.

Subsonic pass: A second FFT pass on a decimated signal at a significantly reduced sample rate, yielding sub-Hz frequency resolution per bin. Examines content below 20 Hz with high precision, detecting DC offset, infrasonic rumble, and subsonic energy.

	High-frequency pass	Subsonic pass
Signal	Native sample rate	Decimated to low sample rate
Bin resolution	Varies by sample rate	Sub-Hz precision
Frequency range	0 - Nyquist	0 - sub-bass region
Purpose	HF extension, DSD curve, brickwall detection	Subsonic energy, DC offset

Table 3 — Dual-FFT architecture summary.

5.2 True Peak — ITU-R BS.1770

True Peak measurement goes beyond simple sample-peak detection. When a digital audio signal is converted to analogue, inter-sample peaks can occur — momentary excursions above the highest sample value that only become visible after reconstruction. Achos measures True Peak using a sinc FIR (Finite Impulse Response) filter with oversampling, as specified by ITU-R BS.1770-4. The FIR kernel applies a linear-phase anti-aliasing filter that models the digital-to-analogue reconstruction process. The result is reported in dBTP.

Interpretation. A True Peak of -0.04 dBTP means the file approaches but does not exceed 0 dBFS after reconstruction. Values above 0 dBTP indicate clipping risk on some DACs. Audiophile masters typically target -1 dBTP or below.

5.3 Integrated Loudness — EBU R128

Integrated Loudness is measured using EBU R128 gating — a two-stage process that excludes silence and very quiet passages. **Stage 1 (absolute gate):** Blocks below -70 LUFS are excluded. **Stage 2 (relative gate):** Blocks more than 10 LU below the Stage 1 mean are further excluded. Achos applies this gating without K-weighting, reporting the unweighted result for maximum analytical transparency.

5.4 Brickwall cutoff detection

A brickwall filter is the tell-tale sign of a sample rate conversion or lossy compression step. Achos detects it using a proprietary slope-based algorithm that measures the rate of energy drop across adjacent spectral bins. When the drop exceeds an internally calibrated threshold over a defined bin window, a cutoff is recorded at that frequency. Threshold-based detection is not used, as it produces false positives on acoustically warm recordings. The slope-based approach is robust to organic acoustic roll-off while remaining sensitive to the sharp cutoffs characteristic of sample rate conversion.

5.5 DSD noise shaping curve

DSD encoding uses a 1-bit delta-sigma modulator that inherently pushes quantisation noise into the ultrasonic frequency range — a process called noise shaping. The result is a characteristic rising noise floor above 20 kHz that reaches its maximum near the Nyquist frequency. This curve is a known and documented property of the DSD format.

Achos verifies DSD authenticity by examining whether this curve is present and follows the expected shape. A file that lacks the curve, or shows a different pattern (such as a drop followed by a rise, indicating PCM-to-DSD conversion), will receive an Inconclusive or Likely Converted verdict. The specific analysis method and calibration thresholds used to identify valid and invalid curve patterns are proprietary to Achos. The curve cannot be synthesised by upsampling a PCM source — it is the definitive DSD fingerprint.

5.6 Genre-aware analysis

Spectral measurements that appear anomalous in one genre may be perfectly normal in another. Achos resolves this using a genre classifier with a four-priority chain: (1) the user's Discogs collection, (2) a Discogs API search, (3) embedded file metadata genre tags, (4) filename parsing. The detected genre informs the AI Filemaster prompt, providing context-appropriate reference ranges.

5.7 Standards reference

Standard	Title	Usage in Achos
ITU-R BS.1770-4	Algorithms to measure audio programme loudness and true-peak level	True Peak measurement (active)
EBU R128	Loudness normalisation and permitted maximum level	Integrated loudness gating (active)
IEC 60908	Audio recording — Compact disc digital audio system	CD reference baseline for upsampling detection
JEITA / JAS	Hi-Resolution Audio definition (96+ kHz / 24-bit)	Quality classification labels
Sony/Philips Scarlet Book	DSD/SACD specification	DSD format validation, noise shaping reference
AES3 / EBU Tech 3250	Digital audio interface standard	Broadcast audio reference
IASA-TC 04	Guidelines on the production and preservation of digital audio objects	Archival quality reference
ISO 266	Acoustics — preferred frequencies	Octave band energy reporting

Table 4 — Standards referenced by Achos. All are freely implementable and do not require royalty payments. Active = used in live measurements.

5.8 AI engine — scope and terms

AI Filemaster and GRaiL use the Google Gemini 2.5 Flash API. Each user must obtain their own API key from Google AI Studio (aistudio.google.com) and configure it in Achos Settings. Achos does not bundle any AI model — all inference occurs on Google's servers. Users must comply with Google's Gemini API terms of service. The specific prompting strategy, system instructions, and output parsing logic used by Achos are proprietary. The AI engine's identity and general methodology are disclosed; implementation details are not.

Part VI — Validation and Reference Testing

6.1 Why self-testing matters

The most reliable way to understand how Achos analyses audio is to test it on files with known, verifiable provenance. A file downloaded directly from a label that records natively to DSD, from a well-documented recording session, with a published technical specification, provides a ground-truth reference.

6.2 NativeDSD free reference files

NativeDSD Music offers a free DSD starter pack containing samples from multiple native DSD recordings across different genres. These files are ideal for reference testing because their provenance is documented and verifiable.

Download: <https://www.nativedsd.com/free-dsd-download/> The pack contains DSD64 (2.8 MHz) samples in DSF format. Download is free and requires registration on the NativeDSD website. Users are encouraged to run Full Analysis and AI Filemaster on these files to establish a personal baseline.

6.3 Expected results — reference files

File	Format	Expected verdict	Key indicator
Barrios — La Catedral (EUDDR1801)	DSD64	Verified	DSD noise shaping curve present. Full Nyquist extension. Wide crest factor.
Vivaldi — Four Seasons, Winter II (40318)	DSD64	Verified	Clean noise curve, full Nyquist extension. Low integrated loudness consistent with classical.
Skidmore — The Impossible Sum (TTK0015)	DSD64 (DXD source)	Likely Converted	Energy drop in HF region before noise shaping rise. DXD-to-DSD conversion — disclosed by NativeDSD.

Table 5 — Reference test results (Achos v1.0.0). The Skidmore file is correctly identified as Likely Converted — NativeDSD documents its DXD origin. The verdict describes the recording chain, not the listening experience.

6.4 Interpreting edge cases

Electronic and synthesised music. Files with no acoustic high-frequency content may show an apparent HF roll-off that resembles upsampling. Genre-aware analysis mitigates this. Inconclusive verdicts are an invitation to verify provenance through the source platform.

DXD-sourced DSD. DXD (352.8 kHz PCM) is a high-quality edit master format commonly used before DSD distribution. Files converted from DXD to DSD will receive Likely Converted. This does not imply poor quality. Check the label's documentation for the recording chain.

Analogue tape transfers. High-quality transfers from analogue masters naturally exhibit elevated noise floors in the higher octave bands. Genre-aware thresholds account for this, but very old recordings with significant tape hiss may produce Inconclusive verdicts.

Appendix A — Glossary

Term	Definition
Bit depth	The number of bits used to encode each audio sample. Higher bit depth = greater dynamic range. 16-bit = CD; 24-bit = hi-res standard.
Brickwall filter	A steep low-pass filter that abruptly removes all frequencies above a set point. Its presence below the Nyquist frequency indicates a sample rate conversion.
Crest factor	The ratio of peak level to RMS level in dB. A high crest factor (>20 dB) indicates wide dynamic range.
dBFS	Decibels relative to Full Scale. 0 dBFS is the maximum digital level. All values are negative.
dBTP	Decibels True Peak. Measures the maximum inter-sample peak after digital-to-analogue reconstruction.
DC offset	A constant non-zero baseline in the audio signal, appearing as energy at 0 Hz. Can cause clicks on playback and is a mastering defect.
DSD	Direct Stream Digital. A 1-bit audio encoding format at 2.8 MHz (DSD64) or higher. Native format of the SACD.
DXD	Digital eXtreme Definition. 352.8 kHz / 32-bit PCM format used as a high-resolution edit master, often before DSD distribution.
EBU R128	European Broadcasting Union loudness standard. Defines integrated loudness measurement with a two-stage gating algorithm.
FFT	Fast Fourier Transform. Mathematical algorithm converting audio from the time domain to the frequency domain.
FIR filter	Finite Impulse Response filter. A linear-phase digital filter used in True Peak measurement to model DAC reconstruction.
Hog Mode	macOS exclusive audio device access. Prevents other applications from using the audio output, enabling bit-perfect playback.
ITU-R BS.1770	International Telecommunication Union recommendation defining True Peak and loudness measurement algorithms.
LUFS	Loudness Units relative to Full Scale. The unit of integrated loudness measurement per EBU R128.
Noise floor	The level of background noise in a recording. Determined by the recording environment, microphone noise, and format bit depth.
Noise shaping	A DSD encoding technique that redistributes quantisation noise into ultrasonic frequencies, producing the characteristic rising high-frequency noise curve.
Nyquist frequency	Half the sample rate. The highest frequency that can be faithfully represented in a digital audio file.
PCM	Pulse-Code Modulation. The standard method of digitally encoding audio by measuring amplitude at regular intervals.
Sample rate	The number of audio samples captured per second, in Hz or kHz. 44.1 kHz = CD; 192 kHz = hi-res PCM.
True Peak	The maximum audio level measured after inter-sample reconstruction, as defined by ITU-R BS.1770.
Upsampling	Converting audio to a higher sample rate without adding genuine high-frequency content. No additional sonic information is created.

Appendix B — Standards and References

All standards listed below are freely implementable and do not require licensing fees or royalty payments.

ITU-R BS.1770-4 (2015)

Algorithms to measure audio programme loudness and true-peak audio level. International Telecommunication Union, Geneva.

EBU R128 (2020)

Loudness normalisation and permitted maximum level of audio signals. European Broadcasting Union, Geneva.

IEC 60908 (1999)

Audio recording — Compact disc digital audio system. International Electrotechnical Commission.

JEITA CP-2307 (2012)

Hi-Resolution Audio definition. Japan Electronics and Information Technology Industries Association.

Sony/Philips Scarlet Book (1999)

Super Audio CD System Description — Physical Specification. Sony Corporation and Philips Electronics.

AES3-2009

AES standard for digital audio — Digital input-output interfacing. Audio Engineering Society.

EBU Tech 3250-E (2004)

Specification of the digital audio interface (the AES/EBU interface). European Broadcasting Union.

IASA-TC 04 (2009)

Guidelines on the production and preservation of digital audio objects. International Association of Sound and Audiovisual Archives.

ISO 266:1997

Acoustics — Preferred frequencies. International Organization for Standardization.

Appendix C — COFI-Based Development Evidence

Human-AI Co-Creation Under the Co-Creative Framework for Interaction Design Evidence Corpus: 49 sessions · 63 git commits · 31 structured prompt documents · 24 in-code diagnostic files Development span: 44 days (2026-04-24 → 2026-06-06) Co-authored by Martín Gallardo Arenas and Claude Sonnet 4.6 (Anthropic) · Santiago, Chile · 2026

C.1 Introduction

This appendix provides formal evidence that the development of Achos Player conforms to the academic framework known as COFI — the Co-Creative Framework for Interaction Design — as defined by Rezwana and Maher (2022, 2023) and published in ACM Transactions on Computer-Human Interaction.

Unlike the typical application of COFI to short creative sessions, the Achos corpus represents a sustained co-creative engagement of 44 consecutive development days. The evidence base is unusually rich: 49 identifiable sessions, 63 git commits co-authored by the human developer and the AI agent, 31 structured prompt documents written by the human before each major session, and 24 in-code diagnostic documents created by the AI agent during active debugging as living analysis artifacts.

C.2 Theoretical Framework

COFI: The Co-Creative Framework for Interaction Design

COFI was developed by Jeba Rezwana and Mary Lou Maher at the University of North Carolina at Charlotte, first presented at ICC3 2021 and published in full in ACM Transactions on Computer-Human Interaction (2023, Vol.†30, No.†5, Article†67). It defines Human-AI co-creativity as a process in which both parties collaborate on a shared creative artifact as partners, with interaction dynamics — including turn-taking, contribution type, and bidirectional communication — as the driving forces of the co-creative process.

COFI Component	Definition
Turn-Taking	The alternating pattern of initiative — who initiates each contribution cycle and how control transitions.
Contribution Type	Function of each contribution: generative (new content), evaluative (assessment), or directive (steering).
Communication Channel	Medium of intent: text instruction, structured specification document, code, visual reference, in-code diagnostic document.
Contribution Similarity	Whether new contribution extends (near) or challenges (far) the partner's prior output.
Social Dynamics	Relational register: collaborative, Socratic, adversarial-constructive, or deferential.
Initiative Level	Degree of autonomy: from fully human-directed to AI-initiated, with mixed-initiative as the productive center.

Table C.1 — Core COFI interaction components (Rezwana & Maher, 2023).

The Docs/ File Pattern — A Novel COFI Communication Channel

The Achos corpus introduces a communication pattern not previously documented in the COFI literature: the pre-session structured specification document. Before each major development session, the human developer wrote a Markdown document — ranging from 175 to 1,220 lines — that specified problem context, desired behavior, design constraints, and explicit rules. These files represent a deliberate formalization of intent before the interactive session began, extending COFI's communication channel component in a novel direction.

C.3 Evidence Corpus Overview

Version / Period	Sessions · Evidence · Primary COFI Dynamic
v0.1.0 (Apr 24) Initial commit	1 session · Git commit · Turn-taking: H establishes working foundation, A receives direction
v0.1.1-v0.1.2 (Apr 28-29) DSD foundation	2 sessions · Git commits · Generative: FIR filter, streaming buffer, DSD metadata pipeline
v0.2.0 (Apr 30-May 3) Dual Mode + Discogs	2 sessions · Git commits · Directive + generative: H defines UI paradigm, A implements dual tab architecture
v0.3.0 (May 12-13) Professional audio path	2 sessions · Git commits · Complex generative: HogMode, BitPerfect, DoP DSD64 — 4 new Swift modules
v0.3.x DSD Bug Sprint ★ Case Study 1	7-9 sessions · 7 commits + 2 stash entries + 5 diagnostic docs · Mixed-initiative iterative debugging
v0.4.0 + DoP Fix ★ Case Study 2 (partial)	3 sessions · 5 commits + 5 diagnostic docs · Root-cause diagnosis: wrong CoreAudio property identified
v0.5.0-v0.5.2 Classification + Spectral	3 sessions · 3 commits · Generative: 9-category engine, dual FFT, ITU-R BS.1770-4 True Peak
v0.6.0 Playlist + Library	4 sessions · 4 commits · Directive + generative: AppKit NSTableView, GRDB SQLite schema, FTSS
v0.7.0 Discogs OAuth + Collection	4 sessions · 3 commits · Generative: OAuth 1.0a PLAINTEXT, collection sync, pressing intelligence
v0.7.3-v0.8.x Sandbox + Keychain	5 sessions · 4 commits · Diagnostic + generative: iterative auth prompt elimination
v0.8.0 AI Filemaster	2 sessions · Docs files (1,220-line UI spec) · Directive: H specifies every visual zone; A implements
v0.9.0 GRail	1 session · 2 commits · 391-line brief · Mixed-initiative: H defines philosophy + persona, A implements 5 parts
v0.9.1 ★ Case Study 3 — 3-Queue Model	4+ sub-sessions · 1 commit · Docs files · Co-negotiation of architectural principle
v0.9.1-v0.9.2 Guardian's Path + Gold system	6 sessions · Docs files · Directive: H specifies dot colors, philosophy, design language; A implements
v1.0.0 ★ Case Study 4 — Verdict Taxonomy	5 sessions · 3 commits · Docs files · Philosophy-driven redesign: epistemic honesty over technical accuracy
v1.0.0 Polish — Onboarding + Menubar	3 sessions · 2 commits · Docs files (367 + 334 lines) · Directive + generative

Table C.2 — Complete corpus session map. ★ = session selected for case study analysis.

Corpus Metric	Value
Total sessions identified	~49
Total git commits	63 (including 2 stash entries)
Docs/ structured prompt files	31 (ranging 175-1,220 lines each)
In-code diagnostic .md files created by AI	24+
Development span	44 days (2026-04-24 → 2026-06-06)
Swift source files created	~80
AI features shipped	3 (Filemaster, GRail, Guardian's Path)
External API integrations	2 (Discogs OAuth 1.0a, Google Gemini)
Git commit co-author	Claude Sonnet 4.6

Table C.3 — Corpus metrics summary.

C.4 Case Studies

Four sessions are analyzed in detail, selected to illustrate the full range of COFI interaction dynamics. Each analysis maps documented evidence to specific COFI components.

Case Study 1 — DSD Decimation & Speed Bug Sprint (Sessions 7-9)

Session IDs	Dates	COFI Type	Artifact produced
Sessions 7, 8, 9 (v0.3.x fixes)	2026-05-16-17	Iterative Mixed-Initiative Debugging (most iterative session in corpus)	Correct DSD playback speed in non-exclusive path + pre-buffer guard + DFF metadata rate fix

After implementing DoP DSD64 playback in v0.3.0, DSD files played at the wrong speed in the non-exclusive audio path. The root cause was a mismatch between the decimation ratio applied in the DSD decoder and the clock rate of the output device — a type of bug characteristically resistant to linear debugging because each subsystem appears internally consistent.

Attempt	Evidence · Action · Outcome
1 — commit 403c54a	Root cause identified. Initial fix applied. Problem persisted — speed still wrong in some configurations.
2-3 — stash d7d11fb, d22e4fa	Two concurrent hypotheses tested. Neither fully resolved. Both stashed — AI agent preserves both for evaluation.
4-5 — commits 311a9c7, 8d6d348	Two sequential attempts to isolate the non-exclusive path. Second refined the first based on real-device test results.
6 — commit 84958b7	Non-exclusive path fully isolated. Speed confirmed correct on real hardware. □
7 — commit 7f5c2f1	New issue discovered: pre-buffer logic firing on PCM tracks. Guard added.
8 — commit 150d854	PCM path validated clean. No startup delay. BitPerfect path confirmed isolated.
9 — commit 94d8a9a	Seek bar showing incorrect duration for DSD files. DFF metadata rate parsing corrected.

COFI Analysis. This session is the clearest example of iterative mixed-initiative co-creation under real-world constraints. The interaction extends across 9 attempts because the hardware environment introduces ground truth that neither party can predict: each fix must be tested on a physical DAC. The two git stash entries represent the AI agent maintaining two parallel hypotheses simultaneously — a non-linear contribution unusual in standard command-response workflows but natural in co-creative debugging. The in-code diagnostic documents created during this sprint represent a novel extension of COFI's communication channel: the AI agent produced living analysis documents inside the codebase as shared working memory.

Case Study 2 — DoP Capability Detection Root Cause (Session 11)

After the DSD speed bug sprint, DSD64 and DSD128 DoP playback via hardware remained non-functional. The most analytically significant finding was the identification of the wrong CoreAudio property in the DoP capability check:

Property	Type	Result
kAudioDevicePropertyClockSources	WRONG	Checks current clock state, not capabilities. Fails spuriously when DoP not active.

kAudioDevicePropertyAvailableNominalSampleRate	CORRECT	Reports what the device can do regardless of current state.
--	----------------	---

COFI Analysis. The identification of the wrong CoreAudio property was not prompted by the human — the human reported behavioral symptoms, and the AI independently traced those symptoms to a specific API misuse. This is an instance of autonomous contribution within a co-creative frame. The five diagnostic documents created in this session served as reference material for subsequent sessions, demonstrating that co-creative output extended beyond the code to include a shared knowledge base.

Case Study 3 — The Three-Queue Architecture (Session 31, Bugs #12-14)

During a stabilization sprint addressing 18 bugs, three threading issues emerged that shared a common root. Rather than addressing each in isolation, the human developer formalized the invariant:

“The audio engine must NEVER wait for the main thread. The main thread must NEVER wait for database or file I/O.”

Queue	Responsibility · Constraint
Audio queue	AVFoundation + FFmpeg operations. Never blocks. Never waits for main or library queue. Pre-existing.
Main queue	SwiftUI rendering only. Never executes I/O. Never waits for audio or library queue. Enforced in this session.
Library queue	Dedicated serial Swift actor for all database reads/writes, FSEvents callbacks, and folder scanning. New in this session.

COFI Analysis. This session represents the most philosophically interesting instance of co-creative emergence in the corpus. The human contribution (the invariant statement) is abstract. The AI agent's contribution (the three-queue model) is concrete. Neither alone constitutes the co-creative output. The three-queue model subsequently governed every new async operation through v1.0.0, demonstrating that co-creative output exceeded the immediate session boundary.

Case Study 4 — Filemaster Verdict Taxonomy Redesign (Session 45)

The Filemaster AI feature had been shipping with accusatory verdict language: .authentic, .suspect, .upsampled, .fraudulent. During real-world testing, the Gayle Skidmore TTK0015 file — a DXD-sourced DSD, fully disclosed by NativeDSD — was correctly identified as non-native DSD but incorrectly labelled "Fraudulent Hi-Res". The human developer identified this not as a classification bug but as a philosophical error.

Old verdict → New verdict	Design rationale
.authentic → .verified "Verified"	No change in meaning. Action: Keep.
.suspect → .inconclusive "Inconclusive"	".suspect" implies guilt without evidence. ".inconclusive" accurately describes what spectral analysis can determine. Action: Review.
.upsampled → .upsampled "Likely Upsampled"	Retained but reworded to "Likely" — spectral inference is probabilistic, not definitive. Action: Consider replacing.
.fraudulent → .converted "Likely Converted"	".fraudulent" imputes intent. ".converted" is descriptive without moral judgment. Does NOT force Replace — a converted file may be high quality and legitimately disclosed. Action: Keep if quality satisfactory.

COFI Analysis. Session 45 demonstrates the most philosophically sophisticated co-creative exchange in the corpus. The human contribution is an ethical position about how AI-generated verdicts should be communicated to users. The AI agent's response — generating a complete taxonomy that operationalizes epistemic honesty — mirrors Case Study 3: the human articulates the "why" and the AI constructs the

"what." The simultaneous update to the Technical Manual extends the co-creative output beyond the codebase, preserving the reasoning chain for future development decisions.

C.5 Cross-Session Analysis

COFI Component Distribution

COFI Component	Presence in Corpus (estimated)
Turn-taking (alternating initiative)	49/49 sessions (100%) — every session involves multiple exchange cycles
Contribution Type: Generative	49/49 (100%) — AI produces code, analysis, documentation in every session
Contribution Type: Evaluative	~40/49 (82%) — H challenges, accepts, or rejects AI proposals
Contribution Type: Directive (H)	~45/49 (92%) — H provides scope, constraints, principles
Contribution Similarity: FAR (challenge/rejection)	~20/49 (41%) — DSD root-cause reframes, taxonomy redesign, threading invariant
Contribution Similarity: NEAR (extension)	~42/49 (86%) — most sessions involve incremental refinement
Mixed Initiative (AI proposes unsolicited)	~15/49 (31%) — root-cause identification, architectural proposals, strategy surfacing
Communication Channel: Structured Docs/ file	31/49 (63%) — unique channel pattern in this corpus
Communication Channel: In-code diagnostic document (AI-authored)	24+ sessions (49%+) — shared working memory pattern
Communication Channel: Git commit (co-authored artifact)	63 commits across 49 sessions — every code artifact jointly signed

Table C.4 — COFI component distribution across the 49-session corpus.

Key Design Decisions as Co-Creative Artifacts

Decision	Origin in co-creative exchange
Three-queue model	H's invariant + AI's architectural operationalization (CS3)
Guardian's Path "no GRail without verified provenance"	H's epistemic principle + AI's gating implementation (Sessions 30, 33)
Slope-based brickwall detection	AI identifies false positive pattern on warm recordings; H approves revised algorithm (Session 14)
Dual FFT architecture	H requests independent subsonic analysis; AI proposes separate decimated-signal pass (Session 14)
Genre-aware thresholds	H identifies spectral bias against electronic music; AI proposes 8-bucket genre system (Sessions 41-43)
Epistemic honesty in verdicts	H triggers redesign with Skidmore case; AI generates epistemically honest taxonomy (CS4)
Gold as AI design language	H specifies visual consistency requirement; AI proposes gold as AI color across all features (Session 37)
DSD bitDepth=1 invariant	AI identifies non-determinism root cause: nil bitDepth → 16 misreports DSD to Gemini (Session 29)
In-memory Keychain cache	AI diagnoses repeated prompt pattern from SwiftUI render-triggered Keychain reads (Session 24)
DoP capability via available sample rates	AI identifies wrong CoreAudio property as root cause (CS2)
Gemini temperature 0.15	AI recommends determinism-focused temperature after non-determinism diagnosis (Session 39)
Pre-flight volume check	AI identifies sandbox auth trigger from unmounted volume resolution attempts (Session 25)

Table C.5 — 12 key design decisions co-created across the corpus. Each represents a conclusion neither party held independently before the session.

C.6 Conclusion

This appendix has demonstrated, through a 44-day corpus of 49 documented sessions and four detailed case studies, that Achos Player was developed through a process conforming to the academic definition of Human-AI Co-Creation as formalized in the COFI framework (Rezwana & Maher, 2022/2023).

1. All six COFI interaction components are present and documentable across the corpus.
2. The interaction pattern is consistently mixed-initiative across all 49 sessions.
3. The Achos corpus introduces a novel COFI communication channel: the structured pre-session specification document (Docs/ file), present in 31 of 49 sessions.
4. Multiple co-creative artifacts — the three-queue model, the epistemic verdict taxonomy, the dual FFT architecture, the DSD bitDepth=1 invariant — cannot be attributed to either party independently.
5. Achos is the first audiophile-grade macOS music player developed under a documented COFI-conformant co-creative process across 44 days, 63 co-authored commits, and 31 structured session documents.

C.7 References

- Lin, Z., Kreminski, M., & Chung, J.J.Y. (2023). Beyond Prompts: Exploring the Design Space of Mixed-Initiative Co-Creativity Systems. arXiv:2305.07465.
- Margarido, S., Roque, L., Machado, P., & Martins, P. (2025). MI-CCy Quantifier: A Framework for Quantifying Mixed-Initiative Co-creativity. Springer EPIA 2025, pp. 3–15.
- Rezwana, J., & Maher, M.L. (2021). COFI: A Framework for Modeling Interaction in Human-AI Co-Creative Systems. ICCV 2021, pp. 444–448.
- Rezwana, J., & Maher, M.L. (2023). Designing Creative AI Partners with COFI. ACM Transactions on Computer-Human Interaction, 30(5), Article 67. DOI: 10.1145/3519026
- Rezwana, J., & Maher, M.L. (2025). Human-Centered AI Communication in Co-Creativity. ACM C&C; 2025. DOI: 10.1145/3698061.3726932
- Zhu, J. et al. (2018). Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. IEEE CIG 2018. DOI: 10.1109/CIG.2018.8490433

Achos Technical Manual — v1.0 · App version: v1.0.0 · Platform: macOS, Apple Silicon Human-AI Co-Creation · AI-Augmented Developer: Martin Gallardo · Coding Agent: Claude AI, Anthropic Developed by the Human-AI Co-Creation team under the direction of Martin Gallardo. © 2026 Achos. All rights reserved.