Tests It Could Fail

Study Protocols for `notepat`, AC Native OS and Typing Cadence

@jeffrey · Aesthetic.Computer · ORCID 0009-0007-4460-4913 · November 2026

Abstract. *Don't Forget the Keyboard* placed `notepat` and AC Native OS in the laptop-performance field and ended on four claims, each paired with the study that would settle it [@jeffrey, 2026c]. This paper writes those studies down before they run: hypotheses, designs, sample sizes, measures, analyses, and the result that would count against each claim. A fifth study covers the typing-cadence feature Aesel added in the same week. The first study, a wider prior-art search, needed only a search engine. I ran it while writing this paper, and it falsified claim 1: children's Scratch pianos already map the letter `c` to the note *C*. The five studies that remain are a 24-hour recall test of the letter-named layout against the DAW's AWSED layout, a cold-boot-to-first-note timing across the fleet, a key-to-sound latency and jitter rig with a tapping task, a listening test of whether a typed melody sounds live, and a test of whether a language model does anything useful with keystroke timing. That last study includes a privacy audit with a decision rule set in advance. Each protocol names its pass line before any data exist.

1. WHY WRITE THE TESTS FIRST

A claim that comes with its own test has to say in advance what result would sink it. *Don't Forget the Keyboard* ended with a table of four claims about Aesthetic.Computer, each with a nearest prior, a status, and a "study that settles it" [@jeffrey, 2026c]. The claims are about systems I built, so I have reasons to read any result kindly. The standard answer to that is pre-registration: fix the hypotheses, the design and the analysis before the data exist, and report any deviation [Nosek et al., 2018, Chambers, 2013]. This paper is that registration, written in the open.

The first protocol was cheap enough to run on the day it was written, and it failed the claim it tested (§2). That is the reason to write the others down. The same kind of check that removed the first claim now applies to the rest.

Table 1 lists the six studies. Each section below gives the question, the hypothesis, the design, the sample, the measures, the analysis, and a *pass line*: the result that the claim needs. Missing the pass line counts against the claim.

2. STUDY 1: IS THE MAPPING NEW?

Question. Has any instrument before `notepat` made the letter on a key play the note with that name, so that `c` plays *C*?

Design. The first paper's search found none in two targeted queries and named where it had not looked: patents, Japanese and German web pianos, and pre-web shareware [@jeffrey, 2026c]. This pass searched those, plus code repositories and Scratch. Each query ran through Aesthetic.Computer's Exa-backed search service, and each hit was opened at its primary page. The kill line was set before searching: *one* instrument that plays at least five of the seven natural notes from their own letters breaks the claim.

Result. The claim failed. Japanese and German web pianos were all positional: the AWSED row, the 61-key number-row layout, or the bottom row [WebMusic, 2026, Tools for Anyone, 2026]. The patents placed notes by how often they are played [US Patent 6,444,888, 2002]. The DOS archive and the code search turned up only positional and ASCII maps. Scratch broke the claim twice. Harvard's ScratchEd site hosts a classroom tutorial in which children make one letter sprite per note: "press the C key and the note will play" [Sienna]. A 2018 project, *Keyboard on Scratch*, plays a whole chromatic octave this way, "press the note letters to play the notes," with sharps on `h`–`l` [per_sam, 2018].

What survives. A narrower fact, not a priority claim. `notepat` carries the letter-named octave into a second octave on `h`–`n` and uses it as a performance instrument and the shell of an operating system. The Scratch result also bears on Study 2. When a child is asked to build a piano out of letters, this is the mapping the child builds. That counts for the learnability hypothesis, but it does not prove it.

3. STUDY 2: LEARNING THE LETTER-NAMED KEYS

Question. Do novices learn and keep key–note pairs faster when the letter on the key is the note's name?

Hypothesis. H2: 24 hours after a short lesson, people who learned `notepat`'s layout recall more key–note pairs than people who learned the AWSED layout, the DAW

#	Claim tested	Design	Sample	Pass line (what the claim needs)	Status
1	No prior maps letters to their own note names.	Multilingual prior-art search: web pianos, patents, code, shareware, Scratch.	7 queries, 1 fetch.	No letter-named instrument found.	**Run. Failed.**
2	Letter-named keys are easier to learn.	Two groups, AWSED or `notepat` layout; recall at 24 h.	40 per group (80% power for $d \geq 0.64$).	`notepat` group recalls more key–note pairs at 24 h, $p < .05$.	Registered.
3	The laptop boots into one instrument.	Power-on to first audible note, by tapper, across the fleet and four setups.	Every fleet model, 10 boots each.	AC Native OS faster than every comparison, on every machine.	Registered.
4	Latency is inside the intimacy budget.	Tapper and two-channel rig; then human tapping to a metronome.	1000 strikes per setup; 24 tappers.	Mean ≤ 10 ms and jitter ≤ 1 ms on HDA; no added tapping variance.	Registered.
5	A typed melody sounds live.	Listening test after Bown et al.; audio-only and with video.	120 listeners online.	Typed excerpts rated no less live than tweaked ones, audio-only.	Registered.
6	Keystroke timing helps a model reply.	Same messages with and without cadence; blind pairwise raters; plus a re-identification audit.	60 messages, 2 models, 5 raters; 20 typists.	Cadence replies preferred, with no timing remarks; audit decides what is sent.	Registered.

Table 1: The six studies. Studies 1–5 test the four claims of @jeffrey [2026c], with claim 1 split into priority (Study 1) and learnability (Study 2); Study 6 tests Aesel's cadence feature (§7). The pass line is fixed here, before any data, and missing it counts against the claim.

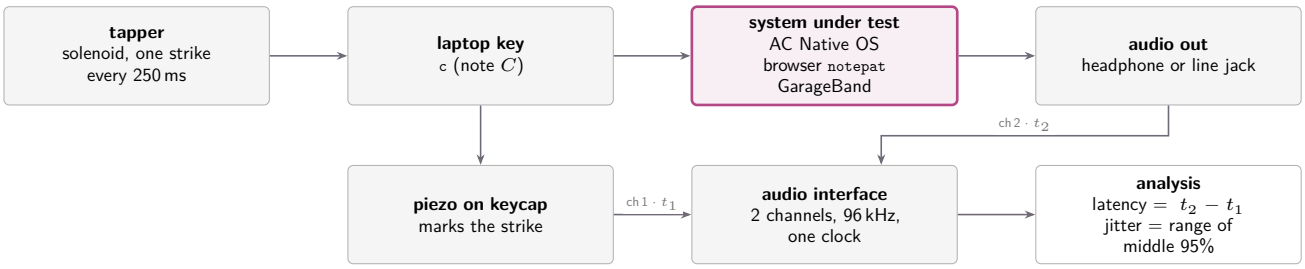


Figure 1: The rig for Studies 3 and 4. The pink box is the system under test. A solenoid strikes the key every 250 ms, as in McPherson et al.'s method [McPherson et al., 2016]. A piezo on the keycap (t_1) and the laptop's audio output (t_2) go to two channels of one interface, so both onsets share one clock. Latency is $t_2 - t_1$. Jitter is the spread of the middle 95% of 1000 strikes.

convention in which `a` plays C and `s` plays D [@jeffrey, 2026b].

Participants. Adults who type but do not play a keyboard instrument. Screening excludes anyone who reports more than one year of piano or keyboard lessons, and anyone who has used `notepat` or a DAW's musical typing. Forty per group gives 80% power to detect a standardized difference of $d = 0.64$ at $\alpha = .05$, two-sided [Cohen, 1988]. Jordà's "low entry fee" predicts a large effect, and this design will not detect a small one [Jordà, 2004].

Procedure. Participants are randomly assigned to one of two layouts, both run in the browser build of `notepat` with key labels hidden. A 15-minute lesson teaches three melodies, written as note names (*E D C D E E E*), as children's songbooks print them. Practice ends when each melody is played once without error, or at 15 minutes. At 24 hours, an online test shows 14 note names in random order (both octaves' naturals) and asks for the key. It then shows one unseen melody as note names and asks for it to be played.

Measures. Primary: the number of correct key–note pairs at 24 h, out of 14. Secondary: lesson time to the first error-free melody; transfer accuracy on the unseen melody; and, as a cost check, the number of piano-position errors. AWSED keeps the piano's white–black geometry and `notepat` does not, so the AWSED group might transfer better to a real piano.

Analysis and pass line. A two-sample comparison of the primary measure (Welch's t; Mann–Whitney if the distribution is badly skewed). Per-item correctness goes into a mixed-effects logistic model with random intercepts for participant and item. H2 passes if the `notepat` group recalls more pairs, $p < .05$. A null result with a confidence interval that excludes $d = 0.64$ counts against the learnability argument made in *Keymaps*.

4. STUDY 3: BOOT TO FIRST NOTE

Question. How long does it take to get from a powered-off laptop to an audible note, and how many human actions does it take?

Hypothesis. H3: on every laptop in the fleet, AC Native OS reaches the first note sooner, and in fewer actions, than every comparison setup on the same machine.

Setups. (a) AC Native OS from USB. (b) The machine's installed operating system, browser, and `notepat` at its URL. (c) GarageBand's Musical Typing, on Macs only. (d) MOD's live USB, measured to its first audible output where the laptop's onboard audio is supported. Each setup is timed on the same machine, so the hardware cancels out.

Procedure. The tapper from Study 4 (Figure 1) strikes `c` every 250 ms from the moment the power button is pressed. A contact microphone marks the power press; the headphone output is recorded. Boot-to-first-note is the time from the press to the first audio onset above −40 dBFS. Steps that need a person (log in, open a browser, type a URL, choose a patch) are done by one operator, as fast as possible, following a written script, and counted. Ten cold boots per setup per machine, with full power-off between boots.

Measures and pass line. Median and range of boot-to-first-note, and the count of human actions. AC Native OS's 7.3 s boot to PID 1 is a lower bound, not the measure [@jeffrey, 2026e]. H3 passes if AC Native OS's median is lower than every comparison's on every machine. A machine where it is not is reported by model, along with the reason.

5. STUDY 4: LATENCY AND JITTER

Question. Is the key-to-sound path inside Wessel and Wright's budget of 10 ms latency and 1 ms variation, and does whatever jitter it has show up in a player's timing [Wessel and Wright, 2002]?

Why jitter. *Where the Microseconds Go* reports a mean of 3–4 ms on HDA codecs and 12–22 ms on SOF, with no jitter figure [@jeffrey, 2026d]. The mean is the wrong number to stop at. Jack, Stockman and McPherson gave percussionists an instrument at 10 ms, at 20 ms, and at 10 ms with ±3 ms jitter. The jittered condition was rated about as badly as 20 ms, while steady 10 ms was rated like zero [Jack et al., 2016]. McPherson, Jack and Moro measured common maker setups and found jitter above 1 ms in every configuration except Bela [McPherson et al., 2016].

Part A: the rig. 1000 strikes per setup, 250 ms apart (Figure 1). Setups: AC Native OS on an HDA laptop and on an SOF laptop; the browser build of `notepat` on the same laptops' installed systems; and GarageBand's Musical Typing on a Mac. Jitter follows McPherson et al.: discard the top and bottom 2.5% of trials, then take the range of the rest [McPherson et al., 2016]. The rig's own error is measured first, with the piezo signal fed to both channels.

Part B: people. 24 participants tap along with a metronome on `f` at intervals of 500, 333 and 250 ms, on AC Native OS and on the slowest setup from Part A, in counterbalanced order. The measure is the standard deviation of asynchrony, the stability index of the tapping literature [Repp, 2005]. Tap times come from the piezo, not from the software, so the system's own jitter cannot hide in the measurement.

Pass line. H4a: on HDA, mean latency ≤ 10 ms and jitter ≤ 1 ms. H4b: asynchrony SD on AC Native OS is no larger than on the slow setup; this is a non-inferiority test with a margin of 2 ms. SOF laptops are expected to fail H4a and are reported as such.

6. STUDY 5: DOES A TYPED MELODY SOUND LIVE?

Question. The liveness critique says a laptop audience cannot see the cause of the sound [Cascone, 2003, Collins, 2003]. Bown, Bell and Parkinson found that listeners, from audio alone, hear *triggering* as more live than tweaking [Bown et al., 2014]. Is a typed melody, which is pure triggering, heard as live?

Design. Following Bown et al.'s approach, listeners rate short excerpts, without knowing how they were made, on how much of what they hear the performer is doing at that moment. There are three sources: (i) a melody typed live on `notepat`; (ii) the same notes played back from a sequence; (iii) a generative patch whose parameters the performer tweaks live. Modality is between subjects: one group hears audio only, the other sees video of the hands and the screen, with the letters visible. 120 listeners online, 60 per group, each rating 18 excerpts in random order.

The trap in the design. In audio only, (i) and (ii) are the same audio. Listeners cannot tell them apart, so any difference between them is noise. That makes (ii) a manipulation check. The claim lives in two other comparisons: typed against tweaked in audio only, and the gain from adding video.

Analysis and pass line. An ordinal mixed model with random effects for listener and excerpt. H5a: in audio only, typed excerpts are rated no less live than tweaked ones; this is a non-inferiority test with a margin of half a scale point. H5b: video raises ratings for typed excerpts more than for tweaked ones (the interaction). If typed excerpts are rated *less* live than tweaked ones in audio only, the argument in the first paper fails.

7. STUDY 6: DOES A MODEL LISTEN TO TYPING?

Question. Aesel now attaches each message's keystroke timing to the prompt it sends to the model. The timing includes interval quantiles, pauses of 2 s or more, bursts, erases and pasted characters, and never which keys [@jeffrey, 2026a]. Does a model reply better with that context, and does sending it leak who is typing?

Part A: usefulness. Sixty real messages from the author's own Aesel sessions, each with its recorded cadence, are sent to two models twice: once with the cadence block and once without. Five raters who do not know which reply is which see the message and two replies in random order, and choose the one that better fits what the person needed. A second set holds the text fixed and swaps in a cadence from a different message: hurried for hesitant, typed for pasted. That set tests whether the replies move with the timing and whether the movement is appropriate. Raters also flag any reply that *mentions* the timing.

Pass line A. H6a: replies with cadence are preferred in more than half of the pairs (binomial test across raters, with message as a random effect), and fewer than 5% of replies remark on the timing. If replies with cadence are not preferred, the feature should be turned off by default.

Part B: the privacy audit. Keystroke timing identifies typists; Killourhy and Maxion's benchmark exists for that purpose [Killourhy and Maxion, 2009]. Twenty consenting typists each produce 30 cadence summaries in ordinary use. A scaled-Manhattan detector, one of the strongest in that benchmark, is trained on each typist's first 15 summaries and tested on the rest [Killourhy and Maxion, 2009]. The equal-error rate is reported twice: once with the raw interval list (`gapsMs`) and once with only the summary statistics.

Decision rule. Set now: if the raw interval list gives an equal-error rate below 15%, Aesel stops sending `gapsMs` by default and keeps only quantiles, pauses and counts. Then the audit runs again. The rule is set in advance so that the result decides the default, not the author's preference.

8. SHARED PROCEDURES

Consent and data. Studies 2, 4B, 5 and 6 involve people. Each participant gets a plain-language description of what is recorded and why. In Study 5 the video shows hands and screens only, never faces. Study 6 records timing and never key identities, and the audit's summaries are deleted after analysis. Where an institution is involved, its ethics review comes first.

Deviations. Anything that changes after this paper is fixed (a sample size, a setup, an exclusion) will be listed in the results paper with its reason. A study that cannot run, such as MOD on unsupported hardware, is reported as not run, not dropped.

Materials. The tapper firmware, the onset-detection script, the lesson melodies, the listening excerpts and the analysis code will go into this paper's source bundle and the `laptop-platter` before data collection starts.

9. LIMITATIONS

These are small studies by one person, and most of their participants will come from the communities around Aesthetic.Computer. Study 2 can detect only a large effect. Study 3 measures the fleet I own, which is surplus hardware, not a sample of laptops. In Study 5, an excerpt chosen by the author is an excerpt the author likes. The excerpts will be cut by a fixed rule (the first 20 s after the first note) to limit that. Study 6's messages are the author's own, so its raters judge replies to one person's typing.

Study 1 is the reminder. A negative search is bounded by where it looked, and the place that broke it, children's coding projects, was not on the first list of places to look.

10. CONCLUSION

The first paper's claims now each have a test written down, and one of them has already failed. The letter-named key is not Aesthetic.Computer's idea; it is what a child builds when asked to make a piano from letters. That leaves a narrower set of claims to defend: a second octave and an instrument built on the mapping, a laptop that boots into it, a latency budget, a liveness argument, and a guess that the gaps between keys say something a model can use. Each of these now has a pass line set before its data exist.

References

Oliver Bown, Renick Bell, and Adam Parkinson. Examining the perception of liveness and activity in laptop music: Listeners' inference about what the performer is doing from the audio alone. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 13–18, London, 2014. doi: 10.5281/zenodo.1178722.

Kim Cascone. Grain, sequence, system: Three levels of reception in

the performance of laptop music. *Contemporary Music Review*, 22 (4):101–104, 2003.

Christopher D. Chambers. Registered reports: A new publishing initiative at Cortex. *Cortex*, 49(3):609–610, 2013. doi: 10.1016/j.cortex.2012.12.016.

Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Lawrence Erlbaum Associates, 2nd edition, 1988.

Nick Collins. Generative music and laptop performance. *Contemporary Music Review*, 22(4):67–79, 2003.

Robert H. Jack, Tony Stockman, and Andrew McPherson. Effect of latency on performer interaction and subjective quality assessment of a digital musical instrument. In *Proceedings of Audio Mostly 2016*, pages 116–123, 2016. doi: 10.1145/2986416.2986428.

@jeffrey. Aesel: `keystroke-timing.mjs`, 2026a. Source, `aesel/src/keystroke-timing.mjs`, Aesthetic.Computer repository.

@jeffrey. Keymaps as social software: Versioned virtual objects via social contract. Aesthetic.Computer papers, 2026b. https://papers.aesthetic.computer/keymaps-social-software-26-arxiv.pdf.

@jeffrey. Don't forget the keyboard: Laptop performance, typed music, and where notepat and AC native OS stand. Aesthetic.Computer papers, 2026c. https://papers.aesthetic.computer/laptop-performance-26-arxiv.pdf.

@jeffrey. Where the microseconds go: Input and audio latency in AC Native OS. Aesthetic.Computer papers, 2026d. https://papers.aesthetic.computer/where-the-microseconds-go-26-arxiv.pdf.

@jeffrey. AC Native OS '26: A bare-metal creative computing operating system. Aesthetic.Computer papers, 2026e. https://papers.aesthetic.computer.

Sergi Jordà. Digital instruments and players: Part I, efficiency and apprenticeship. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Hamamatsu, 2004.

Kevin S. Killourhy and Roy A. Maxion. Comparing anomaly-detection algorithms for keystroke dynamics. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2009. doi: 10.1109/DSN.2009.5270346.

Andrew McPherson, Robert Jack, and Giulio Moro. Action-sound latency: Are our tools fast enough? In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 20–25, Brisbane, 2016. doi: 10.5281/zenodo.3964611.

Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11):2600–2606, 2018. doi: 10.1073/pnas.1708274114.

per_sam. Keyboard on Scratch, 2018. Scratch project 248342290, 26 Sept 2018. "Press the note letters to play the notes." https://scratch.mit.edu/projects/248342290.

Bruno H. Repp. Sensorimotor synchronization: A review of the tapping literature. *Psychonomic Bulletin & Review*, 12(6), 2005. doi: 10.3758/BF03206433.

Sienna. How to make a piano in Scratch. ScratchEd, Harvard Graduate School of Education. Classroom tutorial, undated. https://scratched.gse.harvard.edu/sites/default/files/how_to_make_a_piano_in_scratch.pdf.

Tools for Anyone. Online-klavier, 2026. "Die Tasten werden nach ihrer Position zugeordnet." Retrieved 2026-11-02. https://toolsforanyone.com/de/online-piano/.

US Patent 6,444,888. Musical computer keyboard apparatus and method. US Patent 6,444,888 B1, 2002. https://patents.google.com/patent/US6444888B1/en.

WebMusic. Virtual piano (japanese edition, *baacharu piano*), 2026. Keys `A S D F G H J` play C–B. Retrieved 2026-11-02. https://www.webmusic.app/ja/tools/piano.

David Wessel and Matthew Wright. Problems and prospects for intimate musical control of computers. *Computer Music Journal*, 26 (3):11–22, 2002.

Revision ? · `090f711801` ·
papers.aesthetic.computer/laptop-protocols-26-arxiv

Source bundle

The attached ZIP contains the paper's machine-readable text sources and a plain-text inventory. It supports accessibility, audit, quotation, and reproducible editing; the PDF remains the canonical reading edition. Private evidence, credentials, raw datasets, generated PDFs, and unlisted assets are excluded.