
Designing the Constraints: Eleven Months of Instrument-Making as Art Education

Student name: Roey Tsemah Touitou

Student number: 1099361

Contact: 1099361@hr.nl

Date: June 8th 2026

Educational Institution: Piet Zwart Institute

Program: MA Education in Arts

Table of Contents

Chapter 1: Introduction	05
Chapter 2: Where I Found the Language	13
Chapter 3: How I Studied What I Made	21
Chapter 4: The Learning System	29
Chapter 5: Findings	41
Chapter 6: Discussion	55
Chapter 7: Conclusion	65
References	71

Chapter 1: Introduction

In 2022 I ran a synth DIY workshop in Rotterdam. By the end of a single day, everyone in the room had a working synthesizer in their hands. It was a good day. But building an instrument that someone else has designed is not the same as designing one. Designing asks for something harder: iteration, user testing, documentation, revision, and — the part I find hardest to name — a sense that the work matters enough to be worth finishing. The workshop gave people the satisfaction of a finished object without any of that. They had built; they had not designed.

I kept running into this gap. I founded Synthux Academy, a non-profit, in 2021, and over several years of workshops and online courses I watched the same thing happen again and again: people could build a prototype in a weekend and rarely went any further. The hardware existed. The tutorials existed. The community existed. What was missing was a structure that could hold both support and accountability over months rather than days, and a context that made finishing feel necessary instead of optional.

The residency described in this thesis was my attempt to build that structure. It ran from July 2025 to May 2026. Seven participants were accepted, forming five instrument-design projects. Two of them completed their instruments and presented them at Superbooth 2026, one of the largest electronic instrument fairs in the world. Three did not finish. The question I set out to answer was this:

How does a long-term, hardware-supported, hybrid learning system sustain and develop instrument design competencies in independent practitioners?

That is the question. But this thesis is also about the other side of it — how a system like this distributes participation unevenly. The way I selected participants, the time the program demanded, the technical knowledge it assumed, and the professional framing it took on all created barriers as surely as they created opportunities. Two of the five projects went the whole distance. Three did not, for reasons that range from the clear to the genuinely unknowable. Both of those outcomes are findings, and I try to treat them as such.

How I work, and where the theory came from

I should say something about how I do research, because it shapes everything that follows.

I am a maker, and specifically a designer. I do UX and UI design, I draw faceplates, I design and route circuit boards, and I make videos that teach. The way I work is mostly intuitive and reflective: I build something, I watch it fail or hold, I adjust, and I build again. I did not design this residency from a theory. I designed it the way I design an instrument — with a rough idea of what I wanted it to do, put in front of real people, and changed when reality pushed back. The three-month plan I started with became eleven months. A workshop on visual design appeared because participants needed one. Superbooth was never in the original design; it surfaced in January, when the work needed somewhere to go. None of this was planned in advance and then studied. Running it, and paying close attention while I ran it, is the research.

The theory arrived late. I did not read Donald Schön and then build a reflective residency. I built the residency, and afterwards, looking for language to describe what I had done, I read Schön and recognised myself. The same happened with Etienne Wenger, and with Tim Ingold, whose account of making as a conversation with materials named something I had been watching without words for it. These books did not instruct my practice. They confirmed it, sharpened it, and — this is the part I did not expect — made me feel less alone in it. They showed me that the things I had been doing by instinct had been described before, carefully, by people thinking about how practitioners learn and how things get made.

I want to be honest that this is the reverse of how research is sometimes presented, with theory first and practice as its test. It would be dishonest to pretend I worked that way. And there is something fitting in the order, because Schön's central claim is that practitioners know more than

they can say, and that real reflection happens inside the doing and after it, not before it from a position of theory (Schön, 1983). A thesis in which the making came first and the understanding came after is not a thesis that failed to be rigorous. It is one that did the thing Schön describes. I lived the argument before I read it.

So the theoretical framework in the next chapter should be read in that spirit: not as the foundation I built on, but as the company I found once the building was done. I have kept it small on purpose — three books I have actually thought with, rather than a long list of names. Where they illuminate my work, I say so. Where my context runs past them — because I taught online, asynchronously, across disciplines and time zones, in conditions these books were not written for — I say that too. I'm grateful to my supervisors for encouraging me to write the research this way round — to trust that making first and reading after was a legitimate order, and not one to hide.

A note on AI use: I work with ADHD and through the years learned to build systems that remove the need to rely on my memory. This started with using Notion to document everything — which is why the participant logs and facilitator notes in this thesis are as detailed as they are. AI is an extension of that approach. Throughout this research I used Claude (Anthropic, 2025–2026) as a retrieval tool in three specific ways: to locate passages within my own logs and notes when I knew something was there but not exactly where; to find specific pages in source texts where I had already read and understood the argument; and to challenge my selection of theoretical resources by asking why one book rather than another. In all cases I verified findings against the source material and made the interpretive decisions myself. I also used AI to improve on grammar and check structural coherence across chapters.

The writing, research, practice, data, and ideas are my own. I treat this as consistent with the thesis's own argument: that making and thinking happen together, and that tools shape but do not replace authorship.

Templates and reusable resources can be found in the link below

https://tsemah.notion.site/Designing-the-Constraints-templates-3796331933b880b8a77ade9f7e5f2130?source=copy_link

What this thesis argues

I began by wanting to understand sustained engagement — why some people keep going and others stop. But as I looked back across eleven months of notes, logs, and conversations, a more specific pattern came into focus, and it became the heart of this thesis.

Across the program, participants moved from working inside a technical system someone else had designed toward designing the system themselves. For the first two phases, everyone worked on the same fixed hardware, a small instrument called Spotykach. They wrote their own firmware, invented their own interactions and sounds, but the physical instrument — its knobs, its pads, its shape — was given. In the third phase, that changed completely. Participants designed their own circuit boards. They could now decide which controls existed at all, where they sat, how they looked, and what the instrument communicated before a single note was played.

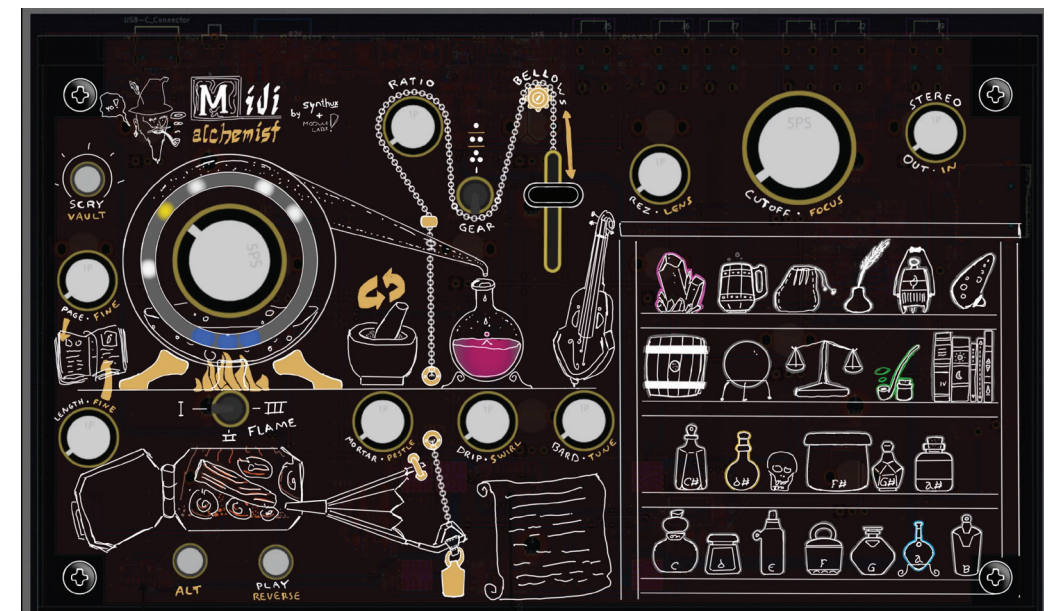
This shift — from designing within constraints to designing the constraints themselves — is the central concern of the thesis. For some participants it opened a real sense of authorship and a new professional identity as instrument designers. For others, the added complexity became the wall they couldn't get over; all three of the projects that didn't finish stalled at or before this transition. I argue that this moment functions as two things at once: a developmental threshold that participants grow through, and a selective mechanism that decides, quietly, who gets to continue. There is a recursion in this I only noticed later — that I was designing the constraints of the residency in much the same way my participants were designing the constraints of their instruments, and learning the same hard lessons about what a constraint includes and what it shuts out.

The instruments were art objects

There is a dimension of this work that the research question doesn't make explicit but that shapes everything else, and I want to state it plainly: the instruments made in this residency were art objects that happened to be playable.

A conventional synthesizer tries to disappear as an object — its interface optimised for function, its look standardised by the industry. The instruments

made here did the opposite. Gerry's MIDI Alchemist carried hand-drawn alchemical imagery across its whole face, coherent enough to stand on its own as a print. Cristian's Sporadic translated the logic of mycelial networks into sound, into animated light, and into a dense drawing of roots and branches — three expressions of a single idea, not a functional device with decoration applied on top. Treating instrument design as art education rather than engineering education meant treating participants as artists developing a practice across technical and aesthetic ground at the same time, and it meant that part of my job as an educator was helping people claim that identity.



Midi Alchemist by Gerry Mayer | Early interface design



Shik S32 | Conventional Midi controller design

This was possible in large part because of who was in the room. The cohort was deliberately mixed. It spanned from fine and audiovisual art — participants thinking in image, performance, and visual meaning — through software and embedded engineering, to hardware and industrial design. Putting these practices in the same sessions produced friction of the most productive kind: an artist asking what an instrument should communicate in a room of engineers asking where to put a control. Some of the most interesting learning in the whole program happened in that meeting between art and engineering, and Chapter 5 treats that crossing as a finding in its own right rather than a happy accident of recruitment.

How the thesis is built, and why it matters

The research follows a qualitative, practice-based methodology, drawing on reflective practice and autoethnography — appropriate to a study in which I am the founder, the hardware designer, the facilitator, and the researcher all at once, deeply embedded in the thing I am studying. The chapters that follow move from the theory I found my language in, through how I studied the residency and how the system was built, to what I found, what it means, and what I would change.

What I hope this offers is threefold: empirical knowledge about how sustained engagement actually works, and fails, in long-term hybrid creative education; a conceptual frame — the constraint-to-authorship transition — for understanding how creative systems scaffold beginners and then ask them to become authors; and a set of practical resources, from Notion templates to hardware files, for anyone trying to build something similar.

The work matters because creative education increasingly happens this way — online, asynchronously, across disciplines, outside institutions. Understanding how these systems hold people, and where they let people fall, is becoming essential for the practitioners who are now learning and teaching inside them. This thesis is one detailed account of one attempt. The value is not in claiming it is the best way, but in describing one way carefully enough that others can learn from it, argue with it, and try something different.

Chapter 2:

Where I Found the Language

As I said in the introduction, I did not begin with theory. I began with a residency I was already running and a feeling that something specific was happening inside it that I couldn't yet name. This chapter is the record of going looking for the names.

I have kept it small on purpose. Rather than survey a field, I have stayed with three books that I read closely and kept returning to: Donald Schön's *The Reflective Practitioner* (1983), Etienne Wenger's *Communities of Practice* (1998), and Tim Ingold's *Making* (2013). Each gave me language for a different part of what the residency turned out to be: how a practitioner learns, how a group of practitioners becomes something together, and how form actually comes into being when someone makes a thing. I treat each one the same way — first for what it lit up in my practice, and then, honestly, for where my situation runs past what the book was written to describe.

Schön: reflection, and the practitioner who sets the problem

Schön describes professional learning as reflection-in-action: practitioners don't apply fixed procedures to clean problems, they think while they work, adjusting as the situation talks back. Beneath that he places knowing-in-action — the competence that lives in the doing and that we mostly can't put into words (Schön, 1983, pp. 49, 68). When I read this I recognised the whole shape of the residency. Participants were never given a method to execute. They met open-ended problems — a piece of firmware that wouldn't compile, a sound that wasn't the sound they heard in their head, a faceplate that needed to mean something —

and they found their way through by trying, watching it fail, and trying again. The biweekly sessions, the Notion logs, the long Discord threads were not separate moments of “reflection” bolted onto the work. They were the reflection, happening inside the work.

It is worth being precise here, because Schön keeps two things distinct that are easy to blur. He describes reflection happening in the midst of action — reflection-in-action, bounded by the window of time in which reflection can still shape what you are doing — and he also describes reflection in retrospect, in the calm of looking back once the work is done (Schön, 1983, p. 60). The residency held both: the live adjusting inside a session, and the looking-back that produced this thesis. What Schön argues against is neither of these but a third thing — the model he names Technical Rationality, in which scientific theory precedes practice and the practitioner’s job is merely to apply it (Schön, 1983, p. 40). When I say the theory in this thesis came after the making rather than before it, this is the distinction I mean: not that reflection only happens afterwards, but that I did not begin from a body of theory and apply it. I began in practice, reflected within it and after it, and went to the books last.

But the concept of Schön’s that matters most to this thesis is his argument that the practitioner’s real skill is not solving the problem but setting it — deciding what the problem even is, and how to frame the situation so that some moves become possible and others fall away. A practitioner, in his account, holds a reflective conversation with the materials of a situation, and that conversation is as much about naming the problem as resolving it (Schön, 1983, p. 40, pp. 94–95).

This is the spine of my central claim. The most important thing that happened in the residency was that participants stopped solving the problem they were handed and started deciding which problem to solve. While they worked on the fixed Spotykach hardware, the problem was given: given this form-factor and controls, what instrument can you make? When they moved to designing their own circuit boards, the problem became theirs to set: which controls should exist at all, where,

looking like what, meaning what? That shift — from designing within constraints to designing the constraints themselves — is exactly Schön’s move from problem-solving to problem-setting, lived out in hardware. I did not have this language while it was happening. I had it afterwards, and it told me the thing I had been watching was real.

Schön also helped me understand a choice I made almost without noticing it. He writes about role frame — how the way a practitioner frames their own role determines what they pay attention to and what counts as a legitimate move (Schön, 1983, p.p 210, 310). The same engineer, framing the job two different ways, sees two different jobs. From the start I framed this residency as art education rather than engineering education, and that frame had teeth. When Cristian couldn’t begin his faceplate, I didn’t read it as a technical problem and send him a tutorial. I read it as an artistic problem — he didn’t yet know what it meant to make a mark for its own sake — and I sent him an artist’s woodcuts instead (Stanley Donwood’s etching works). That was a role-framing decision in Schön’s sense, and it shaped everything that followed for him.

Where Schön bends under my case: almost all of his practitioners are alone in a room, and the reflective conversation he describes is, at heart, a solitary one. Mine was not. The reflection in this residency was distributed — split across a cohort, carried in text, happening asynchronously across several time zones, mostly between people who never shared a physical space until the very end. Reflection-in-action when half the action is a Discord message written at midnight and read the next morning in another country is recognisably Schön’s idea, but stretched into a condition he did not write for. I am not correcting him. I am reporting what his concept looks like when you pull it through a hybrid, networked cohort.

Wenger: becoming someone, and meeting at the boundary

If Schön explains how a single practitioner learns, Wenger explains what happened between my participants. Building on Lave and Wenger's (1991) earlier idea of legitimate peripheral participation — newcomers learning by moving from the edge of a community toward fuller involvement — Wenger's (1998) central insight is that a community of practice is not only a place where knowledge passes around. It is a place where people become certain kinds of people. Identity forms through participation, often without anyone noticing it happening (Wenger, 1998, p.152).

This is the truest description I have of the Gerry–Cristian collaboration. It began as plain technical help and became something else: each man became a different kind of practitioner through the exchange. Cristian, by his own account, stopped thinking of himself only as an engineer and started thinking of himself as a designer — someone who draws, someone who asks what an instrument should communicate. Gerry, as I show in Chapter 5, came by the end to describe Cristian not as someone who had helped him but as a co-creator — itself a piece of identity work. Neither identity was handed out. They emerged from sustained participation, which is exactly what Wenger says identity does.

Wenger also gives me a vocabulary of trajectories — the different ways people move through a community (Wenger, 1998, p.p. 154, 166). Some are inbound, heading toward full membership; some are peripheral and never quite arrive; some lead outward, out of the community altogether. I find I cannot read my own dropout data without this. Participant H never moved from a peripheral position to an inbound one — his prior making experience did not convert into the kind of participation the residency rewarded, and he left early. A and V were deeply engaged and then, somewhere in the move into Phase 3, turned outward and were gone — for reasons I could never recover, as Chapter 5 admits. Reading withdrawal as a trajectory rather than a failure changed how I understood

it, and it is why Chapter 6 treats dropout as data about the system rather than as a verdict on the participant.

But the concept from Wenger that I want to put at the centre, is boundary (Wenger, 1998, p.103-105), and the related ideas of brokering and multimembership. Wenger argues that some of the most important learning does not happen at the comfortable core of a single practice but at the boundary between practices, and that certain people act as brokers, carrying ideas across that boundary, belonging to more than one practice and fully to neither.

This is the frame for the most interesting thing about my cohort. The residency was not one community of practice. It was a small constellation of them rubbing against each other. Gerry came from electrical engineering and industrial design. Cristian, and the pair F and P, came from software and embedded engineering. A and V came from visual and audiovisual art — they thought in images and performance, not in functions. Putting these practices in the same sessions produced friction in the productive sense: an artist asking “what does this communicate?” in a room of engineers asking “where do I put this control?” is a boundary encounter, and the learning happened in that encounter. Cristian the engineer learning to draw, and Gerry moving through three identities the program drew out of him — industrial designer, software practitioner, artist — are not side stories. They are brokering. The diversity of the cohort was not a nice feature of the recruitment; it was the engine of the most distinctive learning the residency produced, and Wenger's boundary is what lets me say so without it sounding like an anecdote.

One writer sits just behind Wenger as a necessary corrective, and I keep the qualification deliberately brief. Contu and Willmott (2003) warn that later readings of community-of-practice theory romanticised participation and looked away from power and exclusion. That caution matters directly to my findings, and I'd extend it with something I watched first-hand: these communities quietly depend on invisible labour — the documenting, mentoring, and onboarding done by a few members

without recognition. The toolchain guide that Gerry started and Cristian extended was exactly this kind of unrecognised, community-sustaining work, and Participant H's early exit may reflect, in part, a community that asked for contribution he was not positioned to give. I use Contu and Willmott not as a pillar but as the honest qualification Wenger's framework needs.

Where Wenger bends under my case: his communities of practice were theorised largely from co-located, physical settings — his central study follows insurance claims processors who share an office. My community was hybrid and mostly online, and the mutual engagement Wenger treats as a given I had to manufacture through Discord, Notion, and biweekly video calls. His framework also tends to assume a single practice with a single notion of competence; my finding lives at the edge he describes as a boundary. Part of what this research offers, then, is a test of whether his ideas hold when the community is networked rather than co-located, and when the boundary between practices is the main event rather than an occasional feature.

Ingold: making as correspondence, not imposition

The third thing I needed language for was the making itself — what was actually happening when a participant sat down to design an instrument. Ingold gave me the cleanest account of it, and it turned out to describe my central finding better than anything else I read.


Ingold's argument in *Making* is built against what he calls the hylomorphic model — the old picture, going back to Aristotle, in which form exists first as a design in the maker's head and is then imposed on inert, waiting matter. He rejects this. Against it he sets making as correspondence: the maker is “a participant amongst a world of active materials,” who joins forces with them in anticipation of what might emerge, rather than standing aloof and “imposing his designs on a world

that is ready and waiting to receive them” (Ingold, 2013, p. 21). Form, in this view, is not transposed from an image onto an object; it emerges through an ongoing back-and-forth between the maker and the material — a process he calls morphogenetic. Making, for Ingold, is a way of thinking, not the execution of thought that happened earlier somewhere else.


This is the truest description I have of what changed in the residency, and it maps almost exactly onto the move I keep returning to. While participants worked on the fixed *Spotykach*, their making was closer to the hylomorphic picture: the form of the instrument was largely given, and the task was to fit ideas into it. When they began designing their own boards and faceplates, the work became correspondence in Ingold's sense — they were no longer fitting a concept into a fixed shape but letting a form emerge through a conversation with their materials. Cristian is the sharpest case. He did not plan the *Sporadic* faceplate and then draw it. He began by tracing, stood back, redrew, followed where the marks led, and the mycelial image emerged through weeks of that back-and-forth. That is morphogenesis, not imposition. He was thinking through the drawing, not illustrating a thought he already had.

What is striking is how closely this sits to Schön. Schön's reflective conversation with the materials of a situation and Ingold's correspondence with active materials are nearly the same insight reached from two different disciplines — one from professional practice, one from anthropology. Both say that the knowing is in the doing and the dialogue, not in a plan that precedes them. The two pillars hold up the same idea, which is part of why the framework feels coherent to me rather than assembled.

Where Ingold bends under my case: he writes mostly about the lone maker in correspondence with wood, stone, or string, and about making in general rather than learning specifically — he is not writing about pedagogy, or about a cohort, or about people teaching each other online. And he does not quite reach the thing I am most interested in:



Spotykach hardware - fixed layout




New hardware, phase 3. Most the yellow area is available for custom design

what happens when the very constraint that scaffolds a beginner later becomes the ceiling an advanced practitioner has to break through — when the fixed instrument that helped you start is the thing you now have to leave behind, at real cost. The movement from imposition toward correspondence, in my case, was not free. It was a threshold some crossed and some couldn't. That dynamic — the constraint-to-authorship transition — is where my contribution sits, and I name it as mine rather than borrowing a citation for it.


The gap this research sits in

These three books gave me a way to understand reflection, identity, and tools. What none of them was written for is the specific situation I found myself describing: a long-term, cross-disciplinary, hybrid learning system running for the better part of a year outside any institution. Most practice-based and community-of-practice research describes either university courses on a semester clock, short maker workshops, or creative communities that already exist and are studied as they are. There is far less on a learning system that is deliberately designed, sustained over many months online, and built across the boundary between art and engineering — and even less that treats its own pedagogical structure as something shaped by power, labour, and exclusion rather than as a neutral support.

So the framework that follows me into the rest of this thesis is not a foundation I built on. It is, as I said, the company I found once the building was done. Schön lets me say that the deepest learning in the residency was participants moving from solving problems to setting them. Wenger lets me say that the cohort's diversity was its engine, that learning happened at the boundary between practices, and that some people brokered across it while others turned outward and left. Ingold lets me say what the making itself became — form emerging through correspondence with materials rather than imposed from a plan. Where my conditions — online, asynchronous, cross-disciplinary, long, and pedagogical rather than solitary — run past what these books describe, that running-past is itself part of what the research has to offer.



Cristian's layout sketches on the collective Miro board




Participants A and V design evolution. On the bottom left the initial design on the fixed hardware, and on the right initial sketch for a custom layout using the new board

Chapter 3: How I Studied What I Made

The question this thesis asks is:

How does a long-term, hardware-supported, hybrid learning system sustain and develop instrument design competencies in independent practitioners?

I have already said where that question came from — from years of watching people build in a weekend and stop, and from wanting to make a structure that could carry them further. This chapter is about how I turned the running of that structure into research, and about the obvious problem that comes with studying something you built, ran, and have a stake in.



Top: Cristian's Sporadic faceplate design
Left: Stanley Donwood's woodcut prints

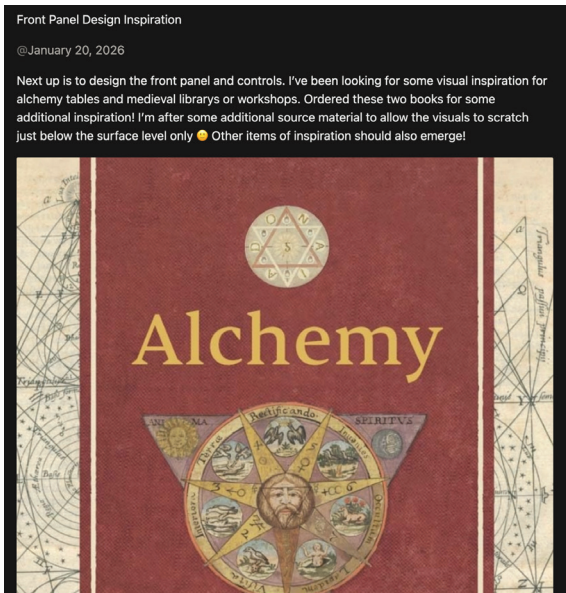


The research is the residency


I did not study the residency from the outside. I studied it by running it, and the running is the data.

This puts the work in the tradition of practice-based research, where the knowledge lives in what gets made rather than only in what gets written about it (Candy & Edmonds, 2018). In this study the made thing is double. There are the instruments the participants built — the obvious artefacts. But there is also the residency itself: its hardware, its tools, its rhythm, its milestones, the decisions I made as I went. The residency is a designed object too, and one of the things I am claiming to have learned about is how to design it. So when I look for findings, I look at both the instruments and the structure that produced them.

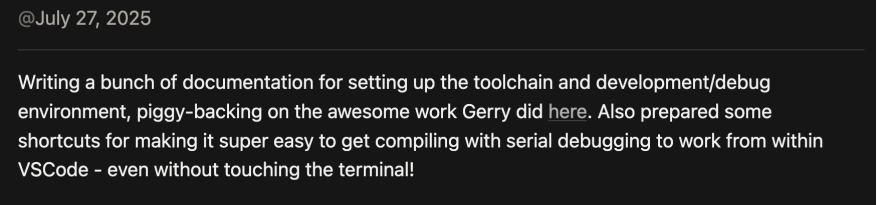
Because I designed that structure responsively rather than in advance, the method was itself reflective practice in Schön’s sense — applied not



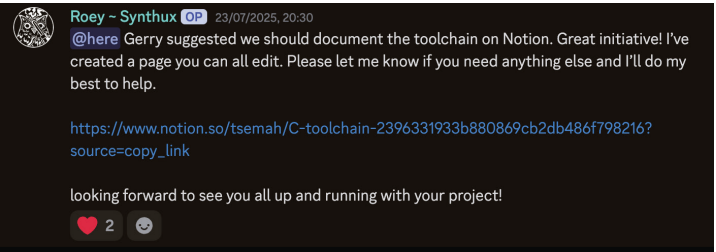
Excerpts from Gerry’s Notion log during the faceplate design phase



to designing a building or treating a client, but to running a residency. I did not lay out eleven months and then observe whether it worked. I laid out a beginning, ran it, watched it talk back, and changed it: a reflective conversation with the situation, carried out over most of a year (Schön, 1983, pp. 94–95). The three-month plan stretched to eleven. Workshops on user experience and visual design appeared when participants hit walls I hadn’t anticipated. The KiCad and PCB tutorials were ran because the work demanded them. Superbooth surfaced in January as a way to



Note from From Cristian’s log regarding the setup guide



Discord note inviting everyone to a dedicated Notion page for the setup guide. Lead by Gerry and Cristian

give the whole thing an ending. In a tidier methodology these changes would be confounding variables to apologise for. Here they are the opposite — they are the evidence. The adjustments are where I can see the system telling me what it needed.

And because I am not a neutral observer of any of this, the work is also autoethnographic — research that uses the researcher’s own embedded experience to understand a practice from inside it (Ellis, Adams & Bochner, 2011). I am Synthux Academy’s founder, the designer of the

hardware, the author of the learning materials, and the residency's facilitator. My session notes, my decisions, my reflections are primary data, not background. This is a real source of knowledge and a real source of bias at the same time, and I deal with that openly later in this chapter rather than pretending it away.

Who was in the room

I announced the residency through the Synthux Academy newsletter — around four thousand subscribers at the time — with a recruitment video describing the hardware, the collaborative aim, and the goal of developing complete instruments over sustained engagement. Thirty-one people applied. I selected seven, on four criteria: the clarity and originality of the concept, evidence of practical making, the reasoning a person gave for why the program fit them, and the diversity of the cohort across skills, backgrounds, and geography. That last criterion turned out to matter more than I understood when I wrote it down.

One pair, participants A and V, I initially declined — their application showed limited technical experience and I wasn't sure. I set them an extra assignment: build your instrument concept in Pure Data and send me a video of it working. They delivered in two days. I was impressed, and I accepted them. I want to flag that this was a judgment applied only to them, on grounds of technical readiness, and it could have been wrong in either direction. I return to it under ethics below.

The seven formed five projects, and the spread of backgrounds is worth setting out plainly, because it became central to what I found. Gerry Mayer, in Chicago, came from electrical engineering and industrial design, and built the MIDI Alchemist, a MIDI-file playback synthesizer. Cristian Pandeale, in Copenhagen, came from embedded engineering and a recent master's in Sound and Music Computing — the most technically experienced of the group — and built Sporadic, a biomimetic audio processor. The pair F and P, in Malmö, came from software



The Synthux Academy Superbooth tent was designed as a playground to encourage visitors to play and explore open source instruments and unconventional unconventional designs



Gerry and Cristian presenting their work at Superbooth 2026

development — F an active contributor to Synthux’s and Monome’s open-source communities, P with fifteen years of professional experience and a background that also included formal study in interactive sound design. Together they built Synk, an oscillator-sync effect. Participant H, in Madrid, came from fine arts and Max/MSP, and was working on a just-intonation drone synthesizer. And participants A and V, in Warsaw, came from visual art and design and from audiovisual art and Arduino, and built M(i/a)crocasm, a polyphonic drone synthesizer. So the cohort ran the whole way from fine and audiovisual art, through software and embedded engineering, to hardware and industrial design. Putting those practices in the same sessions is where some of the most interesting learning happened, and Chapter 5 treats that crossing as a finding rather than a footnote.

On naming: participants signed agreements committing to the residency, but those were not research consent. Two of the completers, Gerry Mayer and Cristian Pandeale, gave explicit informed consent to be named. Everyone else appears under a pseudonym.

The shape of the program — eleven months, biweekly online sessions, Discord and Notion between them, three phases moving from fixed hardware to custom circuit boards — I describe in detail in Chapter 4, since the structure is itself one of the things this research examines. Here I only need to say that the phases are the spine along which everything else in this thesis is arranged.

What I collected

I drew on six kinds of material, gathered as the residency happened rather than reconstructed afterwards.

The first and most continuous were my own session notes — written during or right after each session, recording where participants were, the tone of the room, what was hard, and what I decided to change. These are the primary narrative record, and they doubled as analytical memos: I was

interpreting as I wrote them, not only describing. Alongside them sat the participants’ own project logs in Notion — design notes, videos, firmware, sketches, peer comments. These varied enormously: Cristian’s ran to thousands of words across dozens of entries, Gerry’s was detailed too, others documented sparsely. That unevenness is itself a finding, which I come to in Chapter 5. The third source was the Discord archive, where most of the asynchronous troubleshooting and peer feedback lived; I had started with one channel per project, but it quickly reorganised itself into three — general, technical, and visual design.

The other three brought in voices beyond my own. After Superbooth I ran semi-structured interviews with the two completers, Gerry and Cristian — recorded, transcribed, focused on their experience, motivations, and turning points. The Superbooth presentation in May 2026 which captured how work landed with the wider community. And there was one external check on me: Feedback sessions with my Piet Zwart Institute supervisors at several stages, which I treat as methodological data — a view from outside the residency on what I was finding and how I was interpreting it.

How I made sense of it

The analysis was not a separate stage that began when the residency ended. It was woven through the whole thing, which is what reflective practice means when the practice is also the research: every session note, every change of direction, every debrief was already an interpretation.

After the program closed I re-read everything — session notes, logs, Discord, transcripts — looking for patterns that showed up across more than one source. The themes I trust most are the ones that appeared in several places at once: an interview confirming something a session note had caught, a Discord exchange corroborating an interview, a log documenting what I’d watched happen live.

I have to be honest about an imbalance in this material. My strongest

data comes from the two who finished. Gerry and Cristian gave me long interviews and extensive logs; the three who left appear mostly through my session notes and a few brief exchanges. This is survivorship bias, and it has a direct consequence: my claims about what enabled completion are better grounded than my claims about what prevented it. I try to mark that asymmetry wherever it matters rather than write over it.

The problem of being my own instrument

The obvious objection to this whole study is me. I founded the organisation, I designed the hardware, I wrote the materials, I facilitated the sessions, and I have both a financial stake in Synthux’s success and a personal investment in the idea that this pedagogy works. That is about as much confirmation bias as one researcher can carry.

I have not solved this, but I have tried to work against it in a few concrete ways. I have documented the failures as carefully as the successes — the sessions where engagement fell flat, the dropouts, the fragility of the

whole thing are all in here. I built in outside eyes through my supervisors’ feedback. And I held the participant interviews until after the program had ended, so that people could be critical without it affecting a program still in progress. None of this makes me neutral. It makes my non-neutrality visible, which is the most a study like this can offer.

There is one effect I genuinely can’t account for. The participants knew they were being researched, and their logs were framed as program documentation. Some of them may have written more, or differently, because they knew I would read it. I can note that and not much more.

I’ll add that this embeddedness is also where the value is. As I argued in Chapter 2, reflection for a practitioner isn’t a thing that happens at a clean distance from the work — it happens inside it and just after it. Research and teaching were never fully separable here, and the findings are grounded in sustained, direct contact with the thing being studied. That is the limitation and the strength in the same breath.

Ethics

Three issues deserve to be named directly.

The first is the dual role. As facilitator and researcher I was, in effect, evaluating work I had helped bring into being — an unavoidable tension in practice-based research. I managed it through the external observation and supervisor feedback above, and by treating dropout as a finding rather than a personal failure to be explained away.

The second is power. I selected the participants, set the structure, steered the direction, and could have extended or ended anyone’s involvement. That asymmetry means participants may have held back criticism while the program was live; the post-program interviews were partly an attempt to open a space where honesty cost nothing. It is also why I treat the absence of certain voices — particularly A and V’s, who withdrew in silence — as a limitation of the design rather than a gap I can fill with inference.

Right: Participants A and V interface notes for the fixed Spotykach hardware they wrote alternative firmware for



Left: Participants A and V performing their alternative firmware live

The third is selection itself. The decision to accept or decline was mine alone, against criteria I set alone, and the extra Pure Data assignment I gave A and V applied to no one else. I thought it was a fair test of readiness. It may have been; it may also have been a barrier dressed as fairness. I can't fully settle that, and I'd rather leave the uncertainty on the page than resolve it in my own favour.


Scope and limitations

This is a single case study: five projects, eleven months, one facilitator, one hardware platform, one community. It cannot be generalised. What it can do is give a grounded, detailed account of this one case, and point toward things worth testing elsewhere.


The sample is small by any quantitative measure, which is appropriate to a qualitative, practice-based study that values depth over breadth. The data leans heavily on my own observation and documentation; despite the outside checks, I remain the main instrument of collection, and the risk of selective perception is real. And the program ran almost entirely online — everything except the final days at Superbooth — which means a great deal that happens in shared physical space (the side conversations, the body language, the accidental encounters) simply wasn't available to me. These define the edges of what the research can claim. They don't, I think, undermine what it found.

Chapter 4: The Learning System


The residency was not a set of separate activities. It was a designed system — hardware, tools, a rhythm, a way of putting people together, and a set of interventions — and like everything else in this work it took its final shape by being run rather than by being planned. This chapter describes the system as I built it, and the reasoning behind each part. What the system produced — and where it failed — is Chapter 5. Here I am only laying out the machine.



July 10 Gerry's interaction notes on fixed hardware



November 11 Gerry's initial custom interface design



February 4 Gerry's custom interface design



May 7 Gerry's Superbooth prototype using the latest facaplate design

There were six parts to it: the hardware, the digital infrastructure, the temporal rhythm, the peer architecture, the facilitation, and the public presentation.

The hardware

For the first two phases, from July 2025 through January 2026, everyone worked on the same instrument: Spotykach, a dual-looper running on a Daisy Seed microcontroller. It has a fixed layout — twelve touch-sensitive pads, fifteen programmable analog knobs, CV, MIDI, and audio in and out, and an SD card for memory. Participants didn’t design any of that. They wrote custom firmware that repurposed the hardware they were given.

That was deliberate. Giving everyone the same fixed instrument did two things I wanted. It took the enormous problem of designing a physical interface off the table for a while, so that people could put all their attention into sound and interaction. And it gave the whole cohort a shared language — when five different projects all live on the same board, people can actually help each other, because they are speaking about the same pads and the same knobs. The constraint was there to make collaboration and focus possible.


By late 2025 that same constraint had become a wall. Participants’ concepts were outgrowing the board — Gerry’s instrument needed a different touch interaction and visual feedback which the fixed layout and LEDs couldn’t give; Cristian, participants F, A and V could also benefit from dedicated controls the fixed layout couldn’t spare. So in February

@February 23, 2026

Stared researching vector art tools that would be similar to Procreate for iPad Pro which I’m very familiar with and found Affinity Designer 2 to be a great tool.

I had a crazy idea to see if I could shade the artwork as if it were a medieval wood cut print, to hand calligraphy the letters, and use more medieval accurate branding with an “illuminated” A for in alchemist. I ordered these books to learn more from the land of caligraphy and found that the graphic novel artist Moebius also studied a famous medieval wood cut artist, Albrecht Dürer and let these help get me inspired and up to speed for the task.

Another amazing resource I found was the Plantin-Moretus Museum (Antwerp, Belgium) website which has an INCREDIBLE collection of medieval woodcuts to pour through.
<https://collectie.antwerpen.be/impressedbyplantin/all-woodcuts>



Notes from Gerry’s Notion log

2026 I designed and introduced a new board, built directly from what the residents had been asking for. It kept the microcontroller and circuitry but opened the interface up: participants could now place knobs, buttons, and LEDs wherever they wanted and design their own faceplates in KiCad, the open-source PCB tool. I did this by pushing all the essential circuitry — audio, MIDI, sync — to the top and bottom edges, which freed up better than eighty percent of the board's surface in the middle for participants to fill however they chose.

None of the participants had used KiCad before. So the move from a fixed instrument to a designable one was also a move into a new kind of work: interface design as a visual, industrial, and ergonomic problem, not just a matter of mapping functions onto existing controls. What that shift did to people is the heart of Chapter 5. Here I only want to record that I made it on purpose, late, and in response to what the work was asking for.

The digital infrastructure

Three tools held the residency together between sessions.

Notion was the memory. It ran three databases: the participants' own project logs, my private facilitator session notes, and — from October 2025 — a public firmware database where work could be shared with the wider community. The logs were where participants documented their design as it happened, in whatever mix of writing, video, firmware, and sketches suited them. They could also comment and give feedback to one another. The quality of these logs varied enormously, which turned out to matter.

Discord was the daily conversation — troubleshooting, feedback, and the ordinary social glue of a group that never met in person. I started it with one channel per project, but it ended up changing based on the

**Roey ~ Synthux** OP 16/10/2025, 18:55

Hi [@everyone](#), you might remember that we started a residency program a few months ago... This resulted in some really interesting projects – all **using the Spotykach hardware!** Below you'll find more information and video demos from the creators, some of which you might recognise from our server 😊

We'd love to hear your feedback as we're developing these ideas further. Remember that these are highly experimental firmwares, and **not upgrades to the Spotykach firmware** we've been sharing with you. These are completely different instruments that might turn into an instrument later on, but for now live inside the Spotykach hardware.

 **Yes, it's safe!**
It's completely safe to upload these instruments into your Spotykach. Simply hold the reset button in the back and upload like you normally do.

! Lower your speakers volume
As always, it's good practice to lower the volume of your speakers when trying experimental firmwares.

The instruments:

1. /M(i/a)crocasm - Drone looper meditations by [@Hypovolæmia](#) ☀️ and [@Слепой](#)
2. Sporadic - mycelial networks inspired loops and delays by [@AetherEar](#)
3. Synk - experimental synced oscillator by [@ljudvägg](#) and Philip Linde

More info, videos and firmware files below:
https://www.notion.so/tsemah/Experimental-firmwares-2646331933b880688895f12fd090e19d?source=copy_link

Feel free to share your thoughts below

Discord invitation for the Synthux community to try the alternative firmwares that came from the residency

participants needs into three collective channels — general discussion, technical troubleshooting, and later, visual design. The threading kept conversations branchable while leaving a searchable archive behind, which became its own kind of shared resource.

Miro arrived in Phase 3 as the place to iterate on interface layouts. I built a grid template matching the real dimensions of the new board, so people could sketch placements, move things around, and refine against feedback while staying honest to what the hardware could physically hold.

The rhythm

We met every two weeks, online, for one to two hours on Google Meet. The two-week interval was a deliberate choice, and I tested it against the alternatives in my head before committing. Weekly felt like too much to ask of people doing this alongside full-time jobs and families. Monthly let too much drift in — momentum leaked away between meetings and people arrived having half-forgotten where they were. Two weeks was the compromise that gave people room to work while keeping the thread unbroken. Each session followed a steady shape: ten or fifteen minutes per participant to show where they were, then collaborative problem-solving on whatever was stuck, then facilitation tying individual projects back to the larger questions everyone shared.

Two milestones punctuated that rhythm and gave it somewhere to point. The first was the public firmware sharing on October 16, 2025. The second was Superbooth, proposed in January 2026 as the program's ending. I describe both as components below, and what they did to people's engagement is Chapter 5's business.

The peer architecture

At the start I grouped participants into peer teams, pairing people whose skills complemented each other or whose projects were close enough to share problems. This was where the diversity of the cohort became a design decision rather than just a fact about who applied. I had artists,

software engineers, embedded engineers, and an industrial designer in one virtual space, and I wanted those differences to rub against each other rather than sit in separate lanes.

The formal pairing worked unevenly, and I'm honest about that in Chapter 5. But the most significant peer relationship in the whole residency emerged on its own, outside any structure I built: the collaboration between Gerry and Cristian. It began as plain technical help on Discord in August and September 2025 — Cristian, the more experienced programmer, answering Gerry's questions — and over the months it turned into something genuinely two-way, with Gerry contributing visual and ergonomic design sense, interface feedback, and eventually 3D-printed enclosures for both their instruments. The point I want to make in this chapter is only structural: I built the conditions for this — shared tools, visible work, a reason to talk — but I could not have engineered the relationship itself. The system created the possibility. The participants did the rest.

Facilitation as art education

The most important design decision in the residency wasn't a tool or a schedule. It was the choice to run this as art education rather than engineering education, and that choice defined how I facilitated.

This affected the activities we have taken in each bi-weekly session. The clearest single intervention was the October 16th 2025 user-experience workshop. I spent most of the time in this session with Cristian, whose feature list for Sporadic had become unwieldy and technically framed. The two-hour Miro session introduced design methods: user scenarios, requirement lists, performance-driven thinking. It was the moment I tried to move the thinking from what's possible toward what's meaningful for performance. The session notes from that same day record that Gerry's instrument had also hit its ceiling, needing custom touch interaction and visual feedback which the fixed hardware couldn't provide. Both observations pointed the same way: the Spotykach layout had become a constraint rather than a scaffold.

But the art-educational part of facilitation was usually quieter than a workshop, and it showed up most when a participant's problem looked technical and wasn't. When the time came to design the Sporadic faceplate, Cristian froze in front of his screen, the issue wasn't that he couldn't use the software — it was that he was not used to making a visual decision driven by artistic intention rather than technical need, and didn't know how to begin. The right response wasn't an Affinity Designer tutorial. It was a reference: Stanley Donwood's woodcut prints, whose dense branching forms were close to Cristian's idea of mycelial networks. I offered it not as something to copy but as a way in. I encouraged him to start with tracing the artwork and studying it. Cristian was concerned that the design will be too similar, but as he had to adjust the design to his interface and controls, the traces started taking a shape of their own. In his interview, months later, he described how significant this identity shift was for him. When Gerry's logs started tracing visual references from graphic novel artist Moebius to Albrecht Dürer to a museum's woodcut archive, the facilitation task was the opposite of redirecting him toward efficiency — it was recognising that this art-historical research was legitimate and helping it continue. Offering references, naming an artistic direction, recognising when someone is working through an aesthetic problem rather than a technical one: this is hard to document, but it is the core of what I was actually doing as an educator, and it is where the art-and-engineering boundary in the cohort got worked. Cristian the engineer learning to draw is the sharpest example, and Chapter 5 follows it through.

One more intervention I want to record because I didn't plan it. When participants struggled with the development environment, Gerry wrote a setup guide — and then Cristian extended it, and it grew into a piece of documentation thorough enough that I now use it to onboard future cohorts. It emerged from the cohort rather than from me. Watching that happen was one of the more rewarding moments of the whole program, and it is a good illustration of the system doing something I built the

conditions for but didn't author.

Finally, the Superbooth proposal in January 2026 was itself an intervention, maybe the most consequential one. Proposing that the work be shown at a professional trade fair forced a set of questions the residency had been able to avoid: what does “finished” mean, what is essential versus aspirational, what timeline are we actually committing to. It reframed the work from educational to professional in a single move. What that pressure did — for the people it carried and the people it may have cost — is, again, Chapter 5.

Public presentation

Two moments took the work outside the cohort.

The first was the October 2025 firmware database. Uploading work to a place where anyone in the Synthux community could download and use it — around a hundred Spotykach owners at the time — changed what the work was. Inside the cohort, sharing carries an implicit “help me improve this.” Made public, it becomes “here is something you can actually use,” and that shift demanded different work: manuals clear enough for strangers, tutorials that showed use rather than just features.

The second was Superbooth itself, in Berlin, May 7–9, 2026, where Gerry and Cristian presented their finished instruments alongside commercial manufacturers. This was a far more dramatic threshold than the firmware database — not an educational showcase but a professional one, where the work was met by designers, manufacturers, and musicians as potential products rather than student projects.

Gerry flew all the way from Chicago, USA and Cristian from Copenhagen, Denmark. The booth was designed as a playground to encourage visitors to explore and discover.

A system, not a list

I've described these six parts separately, but they only worked because they held each other up. Miro was useful because the hardware had become flexible enough to design for. The public firmware database worked because there was already an active community on the other side of it. The Superbooth deadline could generate momentum because the biweekly rhythm had built momentum to draw on. Pull one part out and the others lose some of their force. That interdependence is the thing I'm most confident I designed well, and it's the frame I carry into the findings.



Participant F performing his alternative firmware on the Spotykach hardware | Screenshot from live video recording

Chapter 5: Findings

This chapter presents what I found across eleven months of running the residency. Seven patterns emerged from re-reading everything — session notes, project logs, Discord, interviews — and the ones I trust are the ones that showed up in more than one place at once. I present them here as grounded description: what happened, and what each pattern seems to point to. The work of holding these findings up against theory, and of asking which of the uneven outcomes were unjust and which were the legitimate demands of advanced practice, I save for Chapter 6.

A word on what I can and can't see. My evidence is much richer for the two who finished than for the three who left, so my account of what enabled completion is firmer than my account of what prevented it. I try to mark that imbalance where it bites rather than write over it.

The cohort's diversity was the engine

The thing I most underestimated when I designed the residency turned out to be the most generative: how differently the participants thought, depending on where they came from.

The cohort spanned a real range. Gerry came from electrical engineering and industrial design. Cristian, and the pair F and P, came from software and embedded engineering. Participant H came from fine arts. A and V came from visual and audiovisual art. When you put people like this in the same biweekly session, looking at each other's work, they don't just trade techniques — they trade questions and perspectives. The engineers in the room habitually asked where a control should go and what else the firmware could do. The artists asked what the instrument was for, what it should make a person feel, what it should communicate before a note

was played. Neither question is better. But participants kept encountering the other discipline's question as something they hadn't thought to ask, and that encounter is where a lot of the learning lived.

A and V, the most art-led project in the group, performed their polyphonic drone live early on and pushed the whole cohort toward thinking about instruments as things you inhabit rather than operate. The engineers, in turn, gave the artists a grip on what was technically buildable. And the single richest relationship in the residency, between Gerry and Cristian, was itself a crossing of this kind — an industrial designer and an embedded engineer each working in the other's blind spot. I came to think of the boundary between art and engineering not as a gap I had to bridge for people, but as the place where the most interesting work was happening on its own. The diversity wasn't a nice feature of recruitment. It was the engine. In Chapter 6 I give this its proper name — Wenger's idea of learning at the boundary, and of brokering across practices — but the finding stands on its own first: mixing disciplines deliberately produced conversations none of these people would have had alone.

Peer learning became identity work

The clearest example of that crossing is the Gerry-Cristian collaboration, and watching it over eleven months I came to see that it was doing more than transferring skills. It was changing who each of them understood himself to be.

It started as plain help. In July 2025 Gerry was stuck on the C++ toolchain, and the cohort pulled him through. His own log records the moment it turned:

"After much struggling, the team helped me debug the problems I was having with the C++ and VSCode toolchain. I've been able to build the libraries, compile the code, and flash the spotyach with the USB-C cable workflow, success!"

What he did next is the interesting part. Rather than simply move on, he wrote on his to-do list a task: "Document VSCode C++ toolchain setup process in a PPT for future noobs like me." That self-deprecating phrase — "future noobs like me" — quietly turns him from a person who received help into a person who gives it. And Cristian then built on top of Gerry's guide:

"Writing a bunch of documentation for setting up the toolchain and development/debug environment, piggy-backing on the awesome work Gerry did here. Also prepared some shortcuts for making it super easy to get compiling with serial debugging to work from within VSCode — even without touching the terminal!"

The guide grew, became infrastructure, and now onboards future cohorts. But the part I want to record as a finding is the shift in language. By the May 2026 interviews, Gerry no longer described Cristian as someone who helped him. He reached for the word co-creator — unprompted, then adding with a laugh that Cristian might not even want the credit — and said that without Cristian's support on the coding environment, the project might not have continued at all. That re-description is identity work— Gerry positioning the relationship, and himself, differently than he would have at the start.

The cross-pollination ran in both directions. By Phase 3, Gerry had added a fader to his own instrument, directly inspired by Cristian's interface approach. And when Phase 3 introduced KiCad — new software, new workflows, a genuinely unfamiliar process — Cristian stepped up again: during the February review session he recorded the entire tutorial I ran, uploaded and shared it as a reusable guide so every participant could work through the PCB process at their own pace. That is a second piece of community infrastructure he built on behalf of the cohort, and it sits alongside the toolchain guide as evidence of how much the residency depended on labour I had not planned for and could not have produced myself. The structure created the conditions; Cristian decided to use them the way he did. Other peer pairings stayed sporadic. That gap between what the structure enabled and what actually emerged from it is itself a finding.

The hardware scaffolded, then asked for authorship

The single clearest developmental pattern in the data is what happened when the hardware changed under people’s feet.

For the first two phases the fixed Spotykach layout did exactly what I’d hoped: it reduced the decision space, so participants could pour themselves into sound and interaction without also having to invent a physical interface. Then, by mid-October, the same constraint had become a ceiling. My session notes from October 16 catch the moment:

“Gerry’s instrument needs a dedicated approach for touch interaction as well as a visual feedback for file navigation, which is not possible with Spotykach’s limited interface and LEDs. Cristian’s instrument needs dedicated controls — Spotykach’s layout is too big and complex. Everyone’s hitting the constraint ceiling.”

When we took the decision to design custom hardware, the nature of the work changed categorically. Participants stopped working within an interface and started designing one — deciding which controls existed, where, looking like what. Gerry’s February log shows the new kind of thinking, spanning electrical, visual, and ergonomic decisions in a single breath:

“I tightened up some of the controls and did a bit of re-organizing of this front panel sketch to prep for vector art. I wanted to show some of the simulated LED effects, combining both bottom PCB and front panel PCB with side firing LED ideas to outline the objects and leave the touchpad more solid and robust. I decided to rework the bellows to incorporate the knob and 3 position switch more intuitively.”

For some, this was also an artistic threshold, not only a technical one. Gerry, with his design background, moved fluidly into it, and his logs trace serious visual research — from Moebius through Albrecht Dürer to a collection of medieval woodcuts from Plantin-Moretus Museum in Antwerp — that produced the alchemical world covering his faceplate.

What makes Gerry’s development worth naming directly is that none of it came from outside. The industrial designer, the software practitioner, and the artist were all already in him. The program didn’t install any of those identities — it created conditions in which all three had somewhere to go.

Cristian’s development moved differently. He came as an engineer and with no background in visual art. He had never drawn before in any deliberate way. He didn’t know how to begin — not which software to use, but what it meant to make a mark for its own sake. The Donwood reference I describe in Chapter 2 was the way in, and what followed was weeks of drawing that had no precedent in his practice, from traced outlines to a composed faceplate where roots, branches, and network nodes are integrated with the knob and LED placement into a single image. In his interview he described what it had done to him. His confidence had grown in a specific direction: he now saw himself as a designer — an instrument designer, he said, or simply a designer, not just a maker or engineer — and expected to carry that with him.

So the constraint-to-authorship transition enabled real creative agency — and the same transition was where the program lost people. All three dropouts occurred at or before this point. I treat that double edge as the central finding of the thesis, and I interpret it fully in Chapter 6.

Going public changed the work, and a real deadline did what internal ones couldn’t

Two moments took the work outside the cohort, and both changed it.

The October firmware sharing shifted the frame. Work shared inside the cohort carries an implicit “help me improve this”; made downloadable by a hundred-odd strangers, it became “here is something you can use” — and that demanded different work, manuals and tutorials legible to people who hadn’t followed the development. Three projects shared work, and making it legible to outsiders was itself a kind of learning.

Superbooth was the larger threshold, and what it revealed is that external deadlines do something internal ones can't. Internal milestones — a session review, a prototype demo — can be negotiated; if someone isn't ready, the session adapts, at least to a certain degree. The firmware database went live on October 16 2025, and Superbooth happened on May 7-9 2026, whether people were ready or not. That non-negotiability forced the hard questions every maker avoids: what does "finished" mean, what's essential versus aspirational, what's the minimum viable version. In an interview I had with Cristian after Superbooth he described getting renewed energy from the Superbooth idea — the shift from developing alternative firmware for Spotykach to building their own synthesizer and showing it alongside professional manufacturers at a real trade fair.

But the same pressure that carried Gerry and Cristian may have been a wall for others. Whether the professional framing felt enabling or prohibitive to the three projects that didn't finish, I can't say — their accounts aren't in the data. That silence is itself a finding, and I return to it below and in Chapter 6.

The rhythm held, but it was fragile

The biweekly rhythm worked, and it was also the thing most exposed to people's lives.

The two-week interval gave the program a heartbeat — people knew when they'd show work, get feedback, see others move. But rhythm alone couldn't hold motivation when life pressed in. Cristian was candid about this in his interview - he described how the biweekly rhythm sometimes made motivation hard to sustain, especially when external pressures competed for attention. When he started a new job in October, keeping up became much harder. The rhythm didn't adjust — it continued regardless of what was was happening in his life.

What carried motivation across the slack stretches was not the rhythm

but the milestones. When a milestone was clear and close — Superbooth approaching — engagement intensified. In the gap between milestones, in the December-January stretch, it sagged. The finding is that rhythm and stakes do different jobs: the biweekly beat prevented drift, but it was external deadlines that generated urgency, and the structure had no way to bend when a participant's circumstances changed mid-program. That rigidity, I'll argue in Chapter 6, is one of the fixable design faults.

Documentation made some participation visible and some invisible

A quieter pattern ran underneath all of this: the participants who documented heavily became more central, and not always because their work was better.

The logs varied enormously — Cristian's and Gerry's ran to thousands of words across dozens of entries; others left brief notes and long silences. Both are legitimate ways to work; plenty of good making happens without a written trace. But only the documented kind became searchable, citable, reusable infrastructure that the rest of the cohort and future participants could draw on. When Gerry hit a bug, Cristian's written solution was there to find. So documentation did three things at once: it created accountability, it built a shared knowledge base, and it became a place where people positioned themselves as certain kinds of learners.

The cost of this fell unevenly. Writing detailed documentation takes time, confidence in your understanding, and comfort making your thinking public. For someone juggling full-time work or unsure of their technical footing, keeping a rich log was extra labour, not natural reflection — and the same act that made one person's expertise visible could expose another's gaps to the group. The people who became most referenced and most central were, in part, simply the people who could afford to write. I name the ethics of that — invisible labour, and who carries it — in Chapter 6.

The three who left

Three of five projects didn't finish. I treat these not as failures to be explained away but as findings about where the system and people's capacities or circumstances diverged.

Participant H withdrew early, during the initial prototype phase. In an Instagram message in August 2025 he told me he couldn't give the residency the attention he wanted, felt out of his depth with the programming, and hadn't made real progress. What makes his case instructive is that he wasn't a novice maker — he had already built a version of his just-intonation drone synth before joining, using Max/MSP and Synthux hardware. His application proposed porting that instrument to Spotykach's C++ environment and adding generative elements. But the experience he brought was in Max/MSP, a high-level visual environment, and it didn't convert into readiness for what this program quietly required: intermediate C++, and independent debugging. My selection had valued "making experience" without distinguishing between very different technical literacies, and the application had named the time commitment but not the technical prerequisites.

Participants F and P stayed longer. F performed his instrument, Synk, live in October 2025, applying the effect to live audio. P was active through Phase 1 and into Phase 2, leaving in October 2025 when work demands made it impossible. F continued further: after Superbooth was proposed in January he was enthusiastic and actively began collaborating with Cristian, planning to use the same internal PCB and layout for both their instruments so they could split production costs and move faster. It was a concrete and practical response to the new demands of Phase 3. Then,

in February 2026, F sent a message on Discord to say he was leaving — under the combined weight of work, family, and health. Their departures weren't about capability. They were about a structure that couldn't flex when life intensified over eleven months.

A and V are the case I can't explain, and I think the honesty about that matters. Through October 2025 they were among the most engaged people in the cohort — a technically sophisticated polyphonic drone, a live performance, a detailed log. Their Notion log ends on October 15, 2025, but they kept going on Discord and in the biweekly sessions until mid-February 2026. On January 18, after the session where Superbooth was proposed, they answered "Hello! Yeah, we are totally in!" — present for and enthusiastic about the professional presentation. Then, around mid-February, they stopped attending. I emailed them in March to check in and never heard back. Without their account I cannot say whether the Phase 3 demands, the professional framing, their own circumstances, or something else entirely led them out. Their case shows that early engagement doesn't predict completion, and that someone highly visible in one phase can vanish silently in another.

The one thing all three withdrawals share is timing: they cluster at or before the move into Phase 3. That clustering is the strongest signal in the dropout data that the transition from firmware to physical manufacturing — new skills, longer timelines, sustained commitment — introduced barriers the earlier phases didn't have, barriers I had neither clearly communicated nor scaffolded. These aren't individual failures. They're where the system's assumptions — about steady life circumstances, even technical readiness, and the capacity to pick up new skills alone — met a reality that didn't match. Chapter 6 takes up what was fixable in that and what was inherent.

Chapter 6: Discussion

The findings describe a system that worked, unevenly, but productively — it carried two of the five projects through eleven months of hard creative-technical work to finished instruments at Superbooth, and lost the other three at different points along the way. This chapter does two things with that. First, it reads the findings through the three books I found my language in, so that the patterns get explained and not just reported. Second, it asks the harder question those patterns raise: when an outcome is uneven, how do I tell an unjust barrier — a fault in the system I could fix — from legitimate differentiation, the real and defensible demands of advanced practice? Both outcomes, completion and dropout, are findings. The work here is understanding why.

Reframing, and the making of a designer

The central thing the residency produced was a shift in the kind of problem participants were solving, and Schön gives the precise name for it. He argues that the real skill of a practitioner is not solving a given problem but setting it — deciding, in his terms, what to attend to and how to frame the situation so that some moves become possible. For the first two phases, the problem was set for everyone: given Spotykach's fixed controls, what instrument can you make? The move to custom boards handed the problem-setting to the participants. They now decided which controls existed at all, and what the instrument should be. That is Schön's move from problem-solving to problem-setting, made physical in a circuit board (Schön, 1983, p. 40). I had been watching it for months without the word for it.

Schön also explains a decision I made almost without noticing, through his idea of role frame — the way a practitioner frames their own role, which governs what they treat as a legitimate problem in the first place. He puts role frame above problem-setting: framing your role is the deeper, longer-lasting choice, prior to the setting of any particular problem (Schön, 1983, p. 310). I framed this residency as art education rather than engineering education, and that frame decided what I was willing to see as a real problem. When Cristian couldn't begin his faceplate, my frame let me read it as an artistic problem rather than a software one, and so I reached for Donwood's woodcuts rather than a technical tutorial. A different role frame would have produced a different intervention and, very probably, a different Cristian. There's a recursion here I find genuinely useful: I was setting the constraints of the residency in the same way my participants were learning to set the constraints of their instruments — and, like them, discovering only afterwards what my framing had quietly included and shut out.

If Schön explains the shift cognitively, Ingold explains what it was as an act of making. His distinction between the hylomorphic model — form planned in the head and imposed on inert matter — and making as correspondence, where form emerges through a back-and-forth with active materials, names the two halves of the transition precisely (Ingold, 2013, p. 21). On the fixed Spotykach, the form of the instrument was largely given and participants fitted ideas into it; that is closer to imposition. Designing their own boards and faceplates turned the work into correspondence — Cristian not planning the Sporadic faceplate but letting it emerge through weeks of tracing, standing back, and redrawing. He was thinking through the drawing rather than illustrating a finished thought, which is morphogenesis in Ingold's sense rather than the execution of a plan. His interview — quoted in the last chapter, where he calls himself “not only maker or engineer, but designer” — records the

identity that grew out of that change in how he made.

But none of these books quite reaches the thing that makes this finding mine. They describe how a tool or a reframing opens a learner's sense of the possible. None of them dwells on what happens when the very constraint that scaffolds a beginner becomes the ceiling an advanced practitioner has to break through — when the tool that got you started is the thing you now have to leave behind, at real cost. That double edge is where my contribution sits, and it's why all three dropouts cluster at exactly the moment the constraint lifted. The transition that made Gerry and Cristian into designers was the same transition that became a wall for others. Constraint, in other words, is never pedagogically neutral; the decision about when to lift it decides who can keep going.

Learning at the boundary

Chapter 5 argued that the cohort's diversity was the engine. Wenger lets me say why, and the concept that does the work is boundary.

Wenger's communities of practice are usually described as single communities with a shared sense of competence. But he also argues — and this is the part the original version of this thesis underused — that some of the most important learning happens not at the comfortable core of one practice but at the boundary between practices, and that certain people act as brokers who carry meaning across that boundary, belonging to more than one practice and fully to neither. My cohort was not one community. It was a small constellation of them: artists, software engineers, embedded engineers, an industrial designer. The learning I found most striking happened at the seams between them. An artist asking what an instrument communicates, in a room of engineers

asking where to put a control, is a boundary encounter, and Cristian the engineer learning to draw is brokering in its purest form — importing the questions of one practice into another and being changed by the exchange.

This also complicates the older idea of legitimate peripheral participation — Lave and Wenger’s picture of a newcomer moving from the edge of a single community toward its centre. The Gerry-Cristian relationship doesn’t fit that shape, because both men were central and peripheral at once: Cristian central in programming and peripheral in visual design, Gerry the reverse. Neither moved toward a single centre. They met at a boundary and each became a different kind of practitioner through the other. Wenger’s later language — identity forming through participation, people becoming certain kinds of people — fits this far better than a one-directional journey inward. Gerry re-describing Cristian as “co-creator” rather than helper, which I quoted in the last chapter, is identity work in exactly Wenger’s sense.

Wenger’s vocabulary of trajectories also changes how I read the people who left. He distinguishes inbound trajectories (heading toward full membership), peripheral ones (that never quite arrive), and outbound ones (leading out of the community). Participant H never converted a peripheral position into an inbound one — the community asked for a kind of participation his background hadn’t prepared him for, and he left early. A and V look, in retrospect, like an outbound trajectory: deeply engaged and then turning away, for reasons I cannot recover. Reading withdrawal as a trajectory rather than a failure is a way of seeing that people move through a community in more than one direction, and that a program which only recognises the inbound path will mistake every other direction for breakage.

The labour that holds it up

The diversity finding has an underside, and here I lean on the critical voice I keep alongside Wenger rather than as a pillar. Contu and Willmott warn that community-of-practice thinking has often romanticised participation and looked away from power and exclusion. I’d add to that a pattern I watched directly: these communities quietly run on invisible labour — the documenting, mentoring, and onboarding that a few members do without recognition.

Both apply directly to the documentation finding. The toolchain guide that Gerry began and Cristian extended was real, unpaid, community-sustaining work — and it became infrastructure I now reuse to onboard future cohorts. Gerry’s design mentorship and 3D-printed enclosures were the same: gifted labour that served the program’s needs and were repaid only in peer appreciation. The word “mutual” overstates the symmetry. Some participants simply had more to give — more expertise, time, and confidence positioning themselves as teachers — and carried a disproportionate share of the load; others benefited without reciprocating, not from unwillingness but for lack of comparable resources. Participant H’s early exit may reflect more than a technical gap; it may reflect recognising that he could participate mostly as a recipient and not as a contributor, which in a community-driven program can feel untenable.

I don’t think the answer is to pretend this labour doesn’t exist, or to forbid it. Advanced practice often does depend on people contributing beyond their own projects. The honest standard is transparency: saying out loud that the program runs on peer support, valuing more than one kind of contribution — not only technical documentation but emotional and organisational labour, and different ways of working that don’t leave a tidy written trace — and checking, periodically, whether the expectation of contribution is falling hardest on the people least able to carry it.

Going public: accountability, and who it serves

The two public moments created an accountability the internal structure couldn't, and Chapter 5 showed why: external deadlines are non-negotiable in a way session reviews are not. But the same pressure cut differently for different people, and I have to be honest about whom it served, including me.

For Gerry and Cristian, the professional framing was plainly enabling — motivating, validating, a real ending to strive for. For the three who left, I simply don't know; their experience of the professional stakes isn't in the data. That absence is itself a finding. The program never built a way for people to leave with voice — no exit interview, no structured channel for saying “this stopped being for me.” So A and V's withdrawal stays unreadable, and the analysis is missing exactly the perspectives that would test whether Superbooth felt like an opportunity or a wall.

I also need to name an institutional fact plainly. As Synthux Academy's founder, I benefited from the participants presenting at Superbooth: their instruments in our booth demonstrated that the program works and helped our recruitment and legitimacy. I don't think that benefit was illegitimate — but I didn't fully see it at the time, and it sat inside a real ethical tension. Superbooth emerged in January, as I adjusted the program in response to where the work was going, not as a hidden agenda, and the participants genuinely wanted it. But once it existed it became the completion pathway rather than one option among several — and people had joined on a different framing (collaborative learning, hardware support, creative exploration), without being told the program would culminate in a professional trade fair, because that hadn't been decided yet. For some the shift was a gift; for others it may have been a misalignment they couldn't have anticipated. The professional framing isn't the problem; programs that prepare people for professional practice should centre professional contexts. The problem was that it wasn't clear enough at the start for people to choose it, and that there was no dignified second path for those it didn't fit.

What's fixable, and what's inherent

This is where the justice question gets answered directly. All selection is gatekeeping, and gatekeeping is power; I selected seven from thirty-one, against criteria I set alone. The honest task is not to pretend the program excluded no one — every program excludes someone — but to separate the exclusions that reflect fixable faults from those that reflect the real demands of the practice.

Some of what went wrong was fixable, and I own it as design error rather than participant failure. The recruitment never specified the technical prerequisites — intermediate C++, independent debugging — that Phase 3 would demand, so Participant H arrived with real making skills that turned out to be the wrong kind, through no fault of his. The biweekly structure had no give in it; when work, family, or health intensified, as they did for F and P over eleven months, there was no pause mechanism, only attendance or absence. And the goals evolved without being communicated — Superbooth became the destination mid-program while people were still oriented toward the thing they'd signed up for. Each of these is addressable, and Chapter 7 sets out how.

Other limits are not faults but the nature of the thing. Advanced instrument design genuinely requires programming competence and circuit literacy; the question is never whether to require them but how to communicate them and offer routes to build them. With fifteen to twenty minutes per person in a biweekly session, the program could not offer intensive one-to-one scaffolding — that's a real choice about scale and self-direction, defensible as long as it's stated. And any program demanding sustained technical engagement across eleven months will select for a degree of life stability that some capable, committed people simply won't have when their circumstances change. A and V deepen this: they passed a genuine readiness test — the Pure Data assignment, delivered in two days — and outperformed others early, so their withdrawal wasn't inability. Their case marks the limit of what

I can analyse without the dropouts' own voices, and it's the strongest argument for building those voices in next time.

The ethical standard, then, isn't the impossible one of excluding no one. It's transparency: saying honestly what the program requires, who it serves, and what alternatives exist; making selection assess real readiness rather than cultural fit; and accepting that dropout often reveals a limitation in the system rather than a deficiency in the person.

Returning to the research question

The findings answer the research question in five ways, each with a built-in failure mode.

It worked through a peer architecture that generated real identity work — people shaped their practice against each other's different backgrounds and assumptions. But that architecture ran on invisible labour: the ongoing effort of keeping a scattered, multilingual group connected across stalled projects and time-zone gaps. I assumed that would emerge from the group. It didn't. It fell to me, and I hadn't designed for it.

The hardware did what I intended. The Spotykach gave everyone the same starting point; the custom PCB in Phase 3 gave them authorship over the constraints themselves. But the move from one to the other was a threshold, not a gradient. There was no gradual handover. That's where the programme lost people — not through disengagement, but through a sudden increase in demand I hadn't scaffolded.

Public presentation created genuine accountability. Superbooth made the work matter in a way internal deadlines can't. But it also set a single definition of success, and I never built an honest exit for people whose projects weren't going to get there. If you weren't headed to Berlin, there was nowhere else to land.

The eleven-month structure held the programme together. Without the regular rhythm of sessions and checkpoints, the work would have

collapsed into drift. But eleven months assumes eleven months of stable life. Three participants left for reasons that had nothing to do with instrument design, and the structure had no give.

Finally, the tools — Discord, session logs, shared documentation — made certain kinds of participation visible. People who worked through writing, who asked questions in text, who showed up consistently: they're legible in the record. Whether that systematically disadvantaged people who work differently, I can't say. The absent participants left too little trace, and I won't fill that silence with inference.

Taken together, the residency's deepest dynamic was none of these in isolation. It was the movement participants made from working inside a system someone else had designed toward authoring the system themselves — and the uneven distribution of the time, literacy, confidence, and professional orientation that crossing that line requires.

Chapter 7: Conclusion

The residency ran for eleven months and supported seven people across five instrument-design projects. Two of them — Gerry and Cristian — finished, and presented their instruments at Superbooth 2026. Five withdrew, across three projects. I have tried throughout to treat both of those outcomes as findings rather than as a success and a failure, because that is what they are.

The deepest pattern in all of it was not skill acquisition, or peer learning, or finishing, taken on their own. It was the movement participants made from working inside a system someone else had designed toward designing the system themselves — from programming a fixed instrument to authoring a custom one, from solving the problem they were handed to deciding which problem to solve. That transition is the thing this thesis is named for, and I won't re-argue it here; Chapter 6 did that. What I want to do in closing is say what I think the work contributes, be honest about what it can't claim, and point at what I'd want to know next.

What this research contributes

I think the research offers four things:

The first is empirical. There are relatively few close accounts of long-term, self-directed, hybrid creative education running outside an institution — most studies look at semester-length university courses, short maker workshops, or communities that already exist. This is a detailed account of one designed system sustained over eleven months, documented from inside, with its failures recorded as carefully as its successes. The value isn't that it generalises — a single case can't — but that it describes one attempt thickly enough that someone else could learn from it, argue with it, or build differently.

The second is conceptual: the constraint-to-authorship transition as a pedagogical threshold. It names something I haven't found named elsewhere — the way a creative-technical system can scaffold a beginner with a fixed constraint and then, by lifting that constraint, ask them to become an author. It's a lens for seeing both halves at once: how the constraint helps, and how its removal can exclude. It builds on Ingold's account of making as a move from imposing form toward corresponding with materials, but it carries that thinking into territory he doesn't reach: the cost, and the selectivity, of the moment the scaffold comes away, when correspondence becomes a threshold not everyone can cross.

The third is practical, and it matters because this is practitioner research. The residency produced reusable things: the Notion templates, the Discord structure, the toolchain documentation, the open-ended hardware design, and a set of facilitation approaches that treat instrument design as art education rather than engineering. These are artefacts that carry design knowledge into the next program, mine or someone else's.

The fourth is the one I'm least comfortable with and most committed to: a worked example of critical reflexivity. I built the thing, ran it, profited from its success, and then tried to study it honestly — including the places where my selection was a kind of gatekeeping, where the goals shifted under people without being named, and where the program had no dignified way for someone to leave. I don't think practitioner research is worth much if it only documents what worked. This thesis tries to model designing iteratively — responding to the Superbooth opportunity as it arose — while keeping enough distance to see, and say, where those responses created barriers I hadn't intended.

Limitations and absent voices

The honest limits are real. This is one iteration of one program, with one facilitator, one platform, one community, and a self-selected cohort of already-experienced makers. The findings would very likely look different with beginners, a different platform, or participants from different backgrounds.

The limitation I feel most is the absence of the dropouts' own voices, and A and V's especially. They were among the most capable and engaged people in the cohort, they were enthusiastic about Superbooth as late as mid-January, and then they were simply gone, with emails unanswered. I can describe the shape of their withdrawal but not its cause, because the program never built a way for someone to leave with voice — no exit conversation, no channel that made stepping away a legitimate, speakable thing rather than a silence. The result is that the analysis is missing exactly the perspectives that would most test my conclusions. If there is one design lesson I'd insist on for anyone building something like this, it's that.

What I'd carry forward

For anyone designing a hybrid, long-term creative program, the things I'd do differently next time come down to a handful of concrete moves:

- State the technical prerequisites plainly, with honest self-assessment guidance, so people can choose the program knowing what it will ask of them by the end and not only at the start.
- Build temporal flexibility into the structure — pause-and-resume, variable participation — so that a change in someone's work, family, or health doesn't force a binary between full attendance and disappearance.
- Communicate the completion goal clearly, or offer more than one path to finish — a professional showcase for those who want it, a community showcase or portfolio for those who don't — since these ask for different risk tolerance and aspiration.

- Acknowledge and actively support the invisible labour that holds a peer community together, rather than assuming it will emerge for free from the people who can most afford to give it.
- Scaffold the constraint-to-authorship transition instead of letting it arrive abruptly, since that is the precise point where this program lost people.
- Build a way out with voice — exit conversations, anonymous channels, ongoing check-ins that treat disengagement as something a person can name rather than only enact through silence.

For researchers, the gaps I'd point at are long-duration programs, which stay underexamined; dropout read as data about a system rather than as participant failure; the way platforms and tools quietly decide what counts as legitimate participation; and methods that actually capture the people who leave. Three questions in particular seem worth pursuing: whether pause-and-resume mechanisms change completion without killing momentum; whether offering tiered or multiple completion paths retains more people without diluting rigour; and whether the constraint-to-authorship transition shows up the same way in other creative-technical fields, from game design to creative coding.

Closing reflection

The residency succeeded for two projects and didn't for three. The conditions that let Gerry and Cristian finish — technical readiness, stable circumstances, reciprocal mentorship, goals that happened to align with where the program went — were not evenly distributed. Some people lacked them from the start; others lost them along the way.

Designing any educational system means choosing who it serves and under what conditions, and no program serves everyone. The question I've tried to hold honestly is whether a given exclusion reflects an honest alignment with what the work genuinely requires, or a barrier I could have removed with better design. I've argued for transparency as the standard: say what the program asks, who it's for, and what other routes exist; make the criteria assess real readiness rather than people who happen to fit; and accept that not every capable, committed person will finish, and that when they don't, the system is usually telling you something about itself.

This is what practice-based research produces — not proofs, but careful descriptions of what happened when something was tried. I'm not claiming this was the best way, or the only way. I'm claiming I documented one way in enough detail that the next person can take it further, or somewhere else entirely.

The next cohort application was recently announced and I am looking forward to applying my learnings and improving as new challenges arise.

The learning continues.

References

Anthropic. (2025–2026). Claude (claude-sonnet-4-6) [Generatieve AI].
<https://claude.ai>

Candy, L., & Edmonds, E. (2018). Practice-based research in the creative arts: Foundations and futures from the front line. *Leonardo*, 51(1), 63–69. https://doi.org/10.1162/LEON_a_01471

Contu, A., & Willmott, H. (2003). Re-embedding situatedness: The importance of power relations in learning theory. *Organization Science*, 14(3), 283–296.

Ellis, C., Adams, T. E., & Bochner, A. P. (2011). Autoethnography: An overview. *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research*, 12(1), Article 10. <http://nbn-resolving.de/urn:nbn:de:O114-fqs1101108>

Ingold, T. (2013). *Making: Anthropology, archaeology, art and architecture*. Routledge.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.

