Tölvera: Composing With Basal Agencies

Jack Armitage
Intelligent Instruments Lab
University of Iceland
Reykjavík, Iceland
jack@hi.is

Victor Shepardson
Intelligent Instruments Lab
University of Iceland
Reykjavík, Iceland
victorshepardson@hi.is

Thor Magnusson
Intelligent Instruments Lab
University of Iceland
Reykjavík, Iceland
thormagnusson@hi.is

ABSTRACT

Diverse intelligence is an emerging field that views intelligence as a spectrum, drawing insights from e.g. cell-based models and their evolution, and recognising the combined agential properties of biological and engineered materials across disciplines. Within diverse intelligence, basal cognition encapsulates the simpler end of the continuum, focusing on broadly applicable insights from the behaviour of single-celled organisms and simulation. Based on a desire for more diversity of real-time AI in NIME, we developed a library called Tölvera, initially for composable artificial life. In this paper we present Tölvera's design and the practice-based methodology that drove it, reviewing artistic works and emergent themes from design-practice iteration cycles. We reflect on how an early influence of artificial life gave way to an interest in reading Tölvera as a basal art medium, and how unexpected tendencies and capabilities play in perturbative aesthetic tension with compositional decisions. We describe how basal agency research, aesthetics and toolkits are influencing the direction of both design and practice, and review work-in-progress features. Finally, we reflect on how a cell's eye view of agency is shaping our thinking towards AI in NIME.

Author Keywords

Diverse intelligence, basal cognition, agential materiality, artificial intelligence, artificial life, self-organising systems, music performance, practice-based research

CCS Concepts

•Applied computing → Sound and music computing; Performing arts; •Computing methodologies → Intelligent agents;

1. INTRODUCTION

From a broad perspective, we believe that ML increasingly holds a monopoly of thought over what artificial intelligence is, and is often over-generalised to define cognition itself, inducing "distorted and impoverished images of ourselves and our cognition" [40]. In a recent systematic review of ML in NIME [18], Jourdan et al. found a doubling of papers involving machine learning (ML) when comparing 2012-2016 with 2017-2022. This emerging ML monopoly over artificial intelligence (AI) technologies in NIME comes despite ethical and political concerns around the implications for musicians and music culture [29]. Despite this trend, other forms of AI such as artificial life (ALife) [6], robotics [19] and unconventional computing [41], have always too been present at NIME. However, where ML for NIME is becoming more accessible [2] and embeddable [37], tools for working with other kinds of AI are less developed.

We believe that designers of agential and intelligent musical instruments are currently missing out on the advantages that simulation-based approaches like ALife have to offer. Especially in early-stage functional prototyping, when rapid iteration is crucial, ML can be computationally expensive both at training and inference time, often takes large amounts of time, expertise and data to train, and incurs significant action-to-sound latency at inference time, resulting in unmodifiable, non-interpretable, and ethically challenged "black boxes" [18, 37]. In contrast, simulations do not require a dataset, offline training or high-performance compute, can achieve rich behaviour from simple rules, and can be easier to visualise and understand.

Additionally, recent work in ALife is flourishing through new systems like Lenia [8], and hybrid approaches like neural cellular automata (NCA) [27]. Far from being purely experimental and theoretical, it is also being applied in addressing complex challenges such as pandemic apex and social contact prediction [35], and cancer modelling [24]. A framework of diverse intelligence is also emerging [43], which seeks "invariants across evolved, engineered, and hybrid systems to determine what all agents have in common, regardless of their composition". In this context, *in silico* methods are crucial for evaluating models of basic (or basal) cognition [21], that can be compared with open-ended active matter [39].

Motivated by these opportunities, we created *Tölvera*[1]. Initially through artistic practice with ALife-based musical scores, we discovered a gap for a tool that enabled composition of ALife rather than specialising in one system, and one that foregrounded real-time music interaction. In this paper we contextualise Tölvera within the current discourse on diverse intelligence and basal cognition, describe its design and report on early artistic use, and reflect on basal media and diverse intelligence in NIME more broadly.

[1] Website: `https://tolvera.is`
Source: `https://github.com/Intelligent-Instruments-Lab/tolvera`
Discord: `https://discord.gg/ER7tWds9vM`
Install: `pip install tolvera`

NIME'24, 4–6 September, Utrecht, The Netherlands.

Project	Description	Language	OSC	Difficulty	Composability
Lenia	Continuous cellular automata (CCA) based on Conway's Game of Life.	JavaScript	No	Expert	No, single system w/ params
Particle Life	"Vivid structures from rudimentary rules."	Java	No	Beginner	No, single system w/ params
ALIEN	"CUDA-powered artificial life simulation program."	CUDA	No	Expert	Yes, extensive editing features
SwissGL	Minimalist WebGL wrapper for concise low-level simulations.	JavaScript	No	Expert	Yes, low-level GLSL coding
Tölvera	Composable basal agents for music interaction.	Python	Yes	Beginner / Intermediate	Yes, creative coding-style API

Table 1: Comparison of selected popular real-time artificial life tools with Tölvera.

2. BACKGROUND

2.1 Diverse Intelligence & Basal Cognition

Diverse embodiments is a phrase used by Levin [20] and others to refer to "chimeric" assemblages of agential materials that cross traditional disciplinary boundaries:

> "Machine versus organism, evolved versus designed, life versus robotics, and many other distinctions based on prior limitations of technology and imagination should give way to a continuum of chimeric approaches." [10]

We propose that the field of diverse intelligence aligns in certain ways with recent calls for a more-than-human, diffractive approach to NIME design [30], since it seeks to "abandon pre-scientific binary notions of natural kinds (sentient vs mechanical, organism vs machine, etc.) that provide only terminological gatekeeping" [34]. Seifert et al [36] remind us that "the journey from physics to mind is continuous", and Levin encourages us to go "beyond the examples from Earth's phylogenetic tree based on brains to consider and compare agents across the option space of designed and evolved combinations of living, non-living, and software components at all scales" [20]. We can also read efforts to understand aneuronal learning and memory [3, 38] as de-centring the brain and neuronal exceptionalism, drawing a grounded parallel with the broader more-than-human sentiment across NIME and HCI [30].

Searching for common principles that form the *base* level of cognition of diverse embodiments is the aim of research in *basal cognition* [21, 36, 43]. As an example of searching for the simplest possible examples of basal cognition *in silico*, Zhang et al [43] reformulated classic sorting algorithms as one-dimensional cellular automata (CA). Instead of an omniscient top-down algorithm performing the sort, each "cell" was able to swap places with its neighbours. To their surprise, they found "unexpected tendencies and capabilities" in this simple model, such as robustness to perturbation when introducing "broken cells" randomly, and clustering by *algotype* when randomly assigning different sorting algorithms to cells. Exploring such agential materiality [10] through and with music is an exciting prospect that necessitates emphasis on composability. NIME researchers have the opportunity to promote wide-ranging discourse and creative exploration of diverse intelligences, and destabilise the ML monopoly.

2.2 Composability

In the process of discovering requirements for Tölvera, we reviewed currently popular projects (Table 1). We found a distinct contrast between tools that simulate specific systems in detail, such as Lenia [8] and Particle Life, with more generalised tools for designing systems, such as ALIEN and SwissGL. The majority of available code in this space falls into the former category, which cannot easily be recombined, especially by beginners. While tools in the latter category exist, they tend to focus on low-levels of abstraction that require expertise to use. To our knowledge, there are no popular tools that provide means of streaming data in and out, in a way that matches music interaction requirements. In terms of platform choice, web-based tools allow easily sharing visualisations, but servers are required for OSC, and they can not take advantage of the scientific computing and AI ecosystem, which is Python-based. Real-time music interaction in Python is becoming more feasible [12], and projects like Taichi[2] [16] combine high performance with high-level programming, making it increasingly viable for creative AI.

Out of our review of existing tools, *composability* emerged as a key principle, that would address the problems of being limited to exploring single systems, versus having a more expressive framework that only provides low-levels of abstraction. From a programming perspective, Horowitz et al describe composability as giving "the ability to freely combine smaller programmed artifacts into larger ones, to accomplish larger goals" [15]. They contrast this quality with *richness*, "allowing programmers to work with domain-specific visualizations and interactions", and *liveness*, "providing programmers with in-depth feedback about a program's dynamic behavior as the program is edited".

3. TÖLVERA `v0.1.0`

Tölvera is a Python library designed for composing together [15] and interacting with basal [21] agencies [10]. It provides creative coding-style APIs that allow users to combine and compose various built-in behaviours, such as flocking, slime mold growth, and swarming, and also author their own. With built-in support for Open Sound Control (OSC)[3] and interactive machine learning (IML)[4], Tölvera interfaces with and rapidly maps onto existing music software and hardware, striving to be both an accessible and powerful tool for exploring diverse intelligence [20] in artistic contexts.

Inspired by our lab's location in Iceland, the word *Tölvera* is an Icelandic kenning based on *tölva* meaning computer, from *tala* (number) and *völva* (prophetess), and *vera* (being), taken together to mean *number being*.

We have employed Tölvera in various collaborative artistic works, including musical performances, compositions, and multimedia installations[5]. Tölvera's role in these pieces

[2] `https://taichi-lang.org`
[3] `https://github.com/Intelligent-Instruments-Lab/iipyper`
[4] `https://github.com/Intelligent-Instruments-Lab/anguilla`
[5] For peer-reviewed publications, see: `https:`

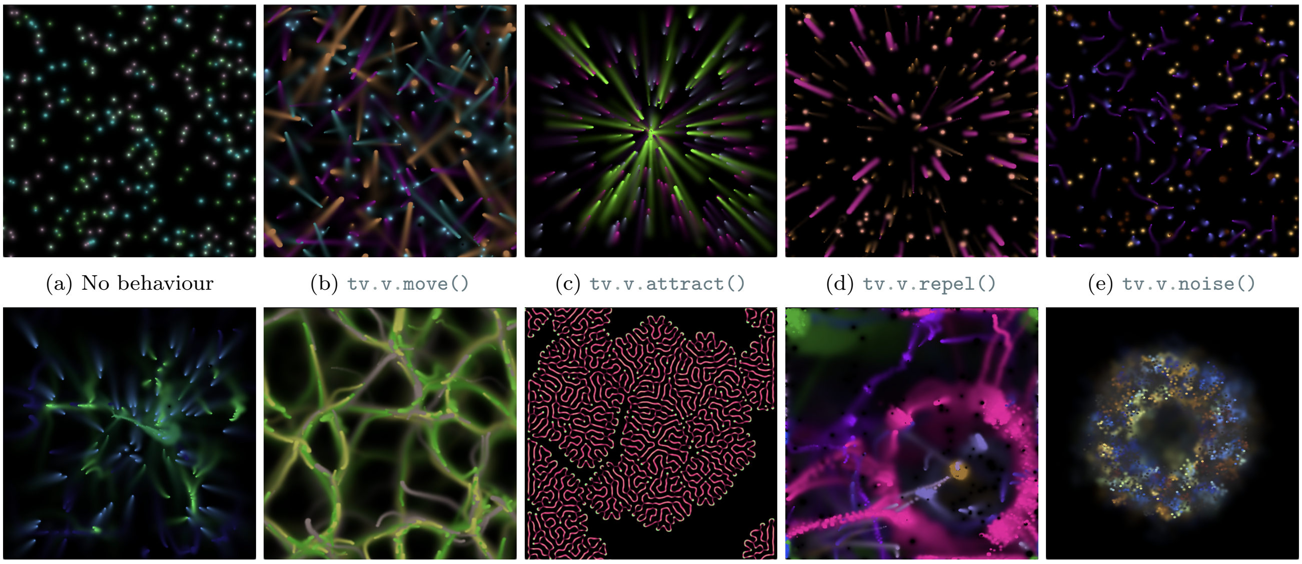

(a) No behaviour (b) `tv.v.move()` (c) `tv.v.attract()` (d) `tv.v.repel()` (e) `tv.v.noise()`

(f) `tv.v.flock()` boids [33] (g) `tv.v.slime()` physarum (h) `tv.v.rd()` reaction-diffusion (i) `tv.v.plife()` particle life (j) `tv.v.swarm()` swarmalators [31]

Figure 1: Examples of behaviours and models available via `tv.v`, some of which are inspired by open source code published by the Taichi community. The top row shows simple, stateless `tv.v` (that do not use `tv.s`), whereas the bottom row are stateful `tv.v`. For example, inter-particle distances are a computational byproduct of `tv.v.flock`, and these values are stored per-particle in `tv.s.flock_p`, and are reusable throughout Tölvera, enabling further composition.

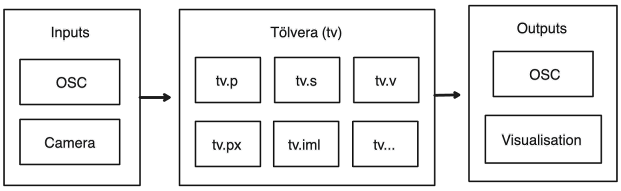

Figure 2: Simplified overview of Tölvera inputs, modules and outputs in `v0.1.0`.

has mainly been "mappable behaviour engine", where interface inputs can control Tölvera programs, and Tölvera runtime data can control interface outputs, in practically any combination (Figure 2). In this way, and to controllable degrees, Tölvera can contribute to the underlying dynamics of a given interactive scenario. It can also add a visual component, and equally has been used without projection in other works. As a creative medium, it has enabled us to practically investigate concepts in diverse intelligence, which has in turn supported our reflections on AI in NIME, which we discuss later in this paper.

Tölvera makes use of Taichi [16], a domain-specific language embedded in Python that enables parallelisation, and is experimental software subject to change. The Tölvera website offers online documentation in the form of an introduction, API guide and reference, code examples and experiments, as well as listings for artworks, publications and a bibliography. In this section we describe the main design features of Tölvera version `0.1.0`, which were developed based on the artistic needs of the pieces described in Section 4.

3.1 `v0.1.0` API

In Python, Tölvera is instanced as `tv`, and its main features are all available via the following sub-objects:

- `tv.p`: Multi-species particle system, where each has a unique relationship with every other species, including itself.
- `tv.s`: Declarative-style global dictionary of n-dimensional (`ndarray`) state structures that can be used by verur, including built-in OSC and IML creation.
- `tv.v`: A collection of behaviours/models including Move, Flock, Slime and Swarm, with more being continuously added. *Verur* can be combined and composed in various ways.
- `tv.px`: Drawing library including various shapes and blend modes, styled similarly to p5.js.
- `tv.ti`: GPU simulation and rendering engine via Taichi. Can be run headless (without graphics).
- `tv.osc`: Open Sound Control (OSC) via `iipyper`[6], including automated export of OSC schemas to JSON, XML, Pure Data (Pd), and Max/MSP.
- `tv.iml`: Declarative-style global dictionary of interactive machine learning instances via `anguilla`[7].
- `tv.cv`: computer vision integration based on OpenCV and Mediapipe[8].

3.2 Multi-species Matrix

Inspired by Particle Life[9], Tölvera has a built-in "multi-species matrix" (Figure 3) such that simulation of complex ecological interactions can be rapidly established, and comprehensively manipulated. Particles (`tv.p`) are assigned "species" identities (represented as an integer), and species

//github.com/Intelligent-Instruments-Lab/tolvera/blob/main/references.bib

[6] https://github.com/Intelligent-Instruments-Lab/iipyper

[7] https://github.com/Intelligent-Instruments-Lab/anguilla

[8] https://developers.google.com/mediapipe

[9] Various implementations exist, for example https://particle-life.com/

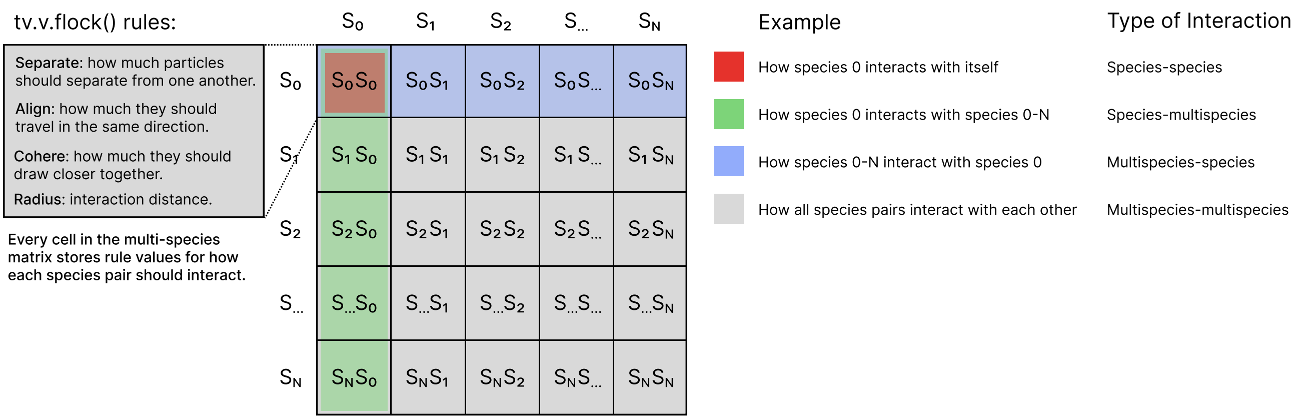

Figure 3: Multi-species matrix (`tv.s.species`) with N species shown on each axis, with example shown based on `tv.v.flock()`'s rules. Every species has a different relationship with each other, including itself, i.e. cell (0,0) shows the 0^{th} species' relationship with itself. As species are implemented as state (`tv.s`), OSC endpoints can be automatically created allowing for dynamically updating rules of individual species pairs, or groups of pairs, or indeed the entire set of rules.

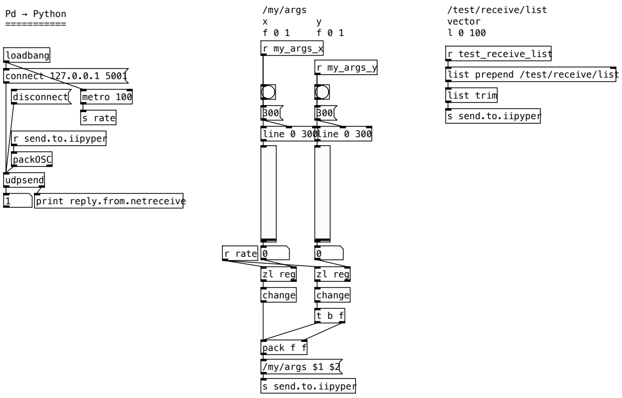

Figure 4: Open Sound Control map (`tv.osc.map`) client generated by Tölvera for Pure Data (Pd) based on the code example in Section 3.4.

can have different relationships with each other, creating a matrix of species-species interactions. These interactions can be manipulated in various ways, and can also be mapped via IML (Section 3.5).

3.3 Composition of Models

The models available via `tv.v` (Figure 1) are decoupled from the particles and can be composed together simply applying them in the same `render` loop. They also have a `weight` parameter to allow further balancing of influence on dynamics. For example, in this listing, the behaviour of the particles will be a compound of the flocking and particle life models, relative to their weight parameters:

```
@tv.render
def _():
  tv.px.clear()
  tv.v.flock(tv.p, 0.8)
  tv.v.plife(tv.p, 0.4)
  tv.px.particles(tv.p, tv.s.species())
  return tv.px
```

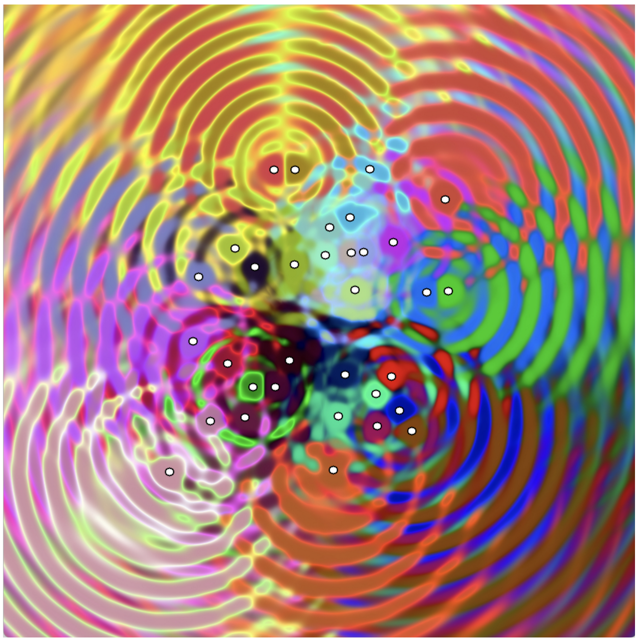

Figure 5: Interactive machine learning (`tv.iml`) real-time map visualisation for input size two (XY axes) and output size three (pixel RGB). Input-output pairs are shown as white circles. This example demonstrates `anguilla`'s ripple interpolator.

3.4 Open Sound Control (`tv.osc`)

`iipyper` integration allows OSC to be written declaratively once in Python, and Tölvera will use this information to export patches for Max/MSP and Pure Data (Figure 4), and JSON and XML schemas.

```
@tv.osc.map.receive_args(
    x=(0.,0.,1.), y=(0.,0.,1.), count=5)
def my_args(x: float, y: float):
  print(f"Receiving args: {x} {y}")

@tv.osc.map.receive_list(
    vector=(0.,0.,1.),
    length=10, count=5)
def my_list(vector: list[float]):
  print(f"Received list: {vector}")
```

3.5 Interactive Machine Learning (`tv.iml`)

`anguilla` integration allows IML mappings to be written declaratively by passing in input and output vector sizes and callback methods, and the maps themselves can be visualised and modified in real-time (Figure 5).

```
tv.iml.flock_p2flock_s = {
  'type': 'fun2fun',
  'size': (tv.s.flock_p.size, tv.s.flock_s.size),
  'io': (tv.s.flock_p.to_vec,
   ↪  tv.s.flock_s.from_vec)
}
```

4. ARTISTIC PRACTICE WITH TÖLVERA

In this section we describe chronologically three examples of artistic practice with Tölvera that shaped its development (Table 2 & Figure 6), aiming to offer empirical insights following Gioti et al's reportage on artistic practice with ML [13]. Unlike approaches where collaboration occurs after development, we embraced early artistic contextualisation to discover what the most urgent practical necessities were, and iterate from there. Quotes from collaborators in these sections stem from conversations with the first author that took place after the performances.

4.1 MOTHERBIRD (May 2023)

MOTHERBIRD (Figure 6a) is a trio for augmented flute (Jessica Shand), electronics (Manuel Cherep) and artificial life (Jack Armitage). It was premiered at TENOR 2023 conference, Northeastern University, Boston, MA, USA and a video archive can be found online[10]. The piece "reimagines the centuries-old flute-as-bird archetype in a 21st-century context in which anthropogenic climate change has drastically altered the soundscapes of the natural world" (programme note). For this performance, we trained two RAVE models; one on Shand's flute recordings, and one on a mixture of recorded and synthetic birdsong[11] (synthesised by Cherep [9]). Cherep operated a Max/MSP patch that controlled the RAVE models, and sent flute audio features to a Tölvera program controlled by TA that mimicked murmuration patterns using `tv.v.flock`, which in return sent flock features back to Max/MSP.

Shand found the murmuration "mesmerizing to watch, and keeps me on my toes as a performer", noticing "the various scales of interactions among the flocking Boids, from the micro to the macro level". As a creative coder, they noted the need for "balance between control and guidance by the system" to "make the interaction even tighter". Notably, Cherep extended the Pixels module with mesh winding to create the MOTHERBIRD character, commenting that Taichi is "a bit more difficult to debug, but it suddenly makes some applications possible in real time [in Python]".

4.2 FerroNeural (May 2023)

FerroNeural (Figure 6b) is a duo for magnetic discs (Nicola Privato) and artificial life (Jack Armitage). It was premiered at Instruments, Interfaces, Infrastructures conference 2023[12], Harvard University, Boston, MA, USA and a video archive can be found online[13]. The title "is a portmanteau referencing to ferromagnetism and artificial intelligence as the protagonist of this performance, whose intertwining agencies give rise to a rich visual and sonic experience" (programme note). In the piece, Tölvera is projected top-down, with two attractors calibrated to two magnets on a table, which Privato's magnetic disc controllers [32] interact with to create a synchrony between sound, gesture and visual. The Tölvera behaviour is a composite of `tv.v.flock` and `tv.v.slime`, creating a "slime flock". Control of RAVE neural audio synthesis models [7] is shared between Tölvera and the discs. The monochrome visual presentation was in reference to musical notation, exploring the concept of agential scores. This was the first piece also drove the need for OSC mapping and client patch generation features (Section 3.4) to facilitate collaboration. It was also using a single species, with state managed manually, which was labour intensive. Privato reflected on the piece that "I would love to lose part of that control [over the mapping], and maybe let the beings reconfigure mapping, sounds and gestures. What interests me is finding and exploring the fine line where I actively negotiate my agency with that of the system, and where none of the two ultimately prevails."

4.3 Pandora's Mycophony (August 2023)

Pandora's Mycophony (Figure 6c) is a duo for two live coders (Celeste Betancur and Jack Armitage). It was premiered at AIMC in June 2023 [4], but here we discuss a second iteration performed at Mengi, Reykjavik, Iceland in August 2023. The piece is an "audiovisual dreamscape" and a "re-imagining of Pandora's story, where the contents of her jar are bioluminescent swarming spores" (programme note). Betancur was live coding a new system called Pandora's Dream [5], which integrated new AI features of the ChucK language [42], and the CHmUsiCK library[14]. TA live coded OSC parameters in Tölvera via TidalCycles [26]. Tölvera was projected on to a custom holographic display, with the two performers looking through it at each other, exploring eye contact in live coding performance. This piece used multi-species `tv.v.flock`-`tv.v.slime` composite interactions and two IML instances, one from Pandora to Tölvera, and vice versa, creating a feedback loop that was central to the composition. Betancur reflected that IML worked "as an initial experiment in bridging the two platforms", since it "easily remaps inputs and outputs in a manner that enhances flexibility and versatility [...] this ap-

[10] https://youtu.be/kuxYIYgPrTs?t=2705

[11] https://huggingface.co/Intelligent-Instruments-Lab/rave-models

[12] https://sites.harvard.edu/instruments-interfaces-infrastructures/

[13] https://youtu.be/Ej7XJBZ6GLE

[14] https://github.com/celestebetancur/CHmUsiCK

	Piece	Month ('23)	Ensemble	Instruments	Species	Composition	Mapping	Projection
4.1	MOTHERBIRD	May	Trio	Flute & RAVE	Single	`tv.v.flock`	Manual	Wall
4.2	FerroNeural	May	Duo	DMI & RAVE	Single	`tv.v.flock` & `tv.v.slime`	Manual	Table
4.3	Pandora's Mycophony	Aug	Duo	ChucK	Multi	`tv.v.flock` & `tv.v.slime`	IML	Hologram

Table 2: Summary of three audiovisual performances with Tölvera.

(a) *MOTHERBIRD*, trio with flute & electronics, May 2023.

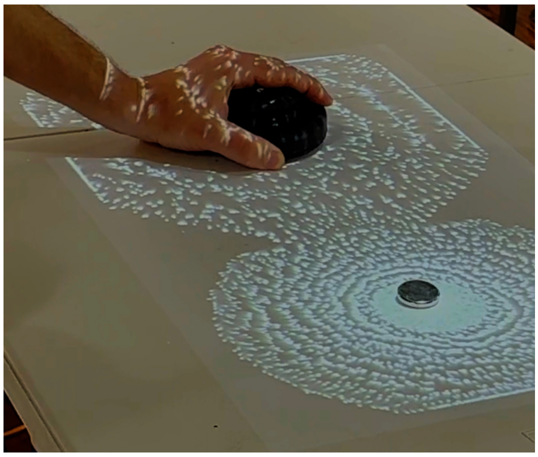

(b) *FerroNeural*, duo with magnetic discs, May 2023.

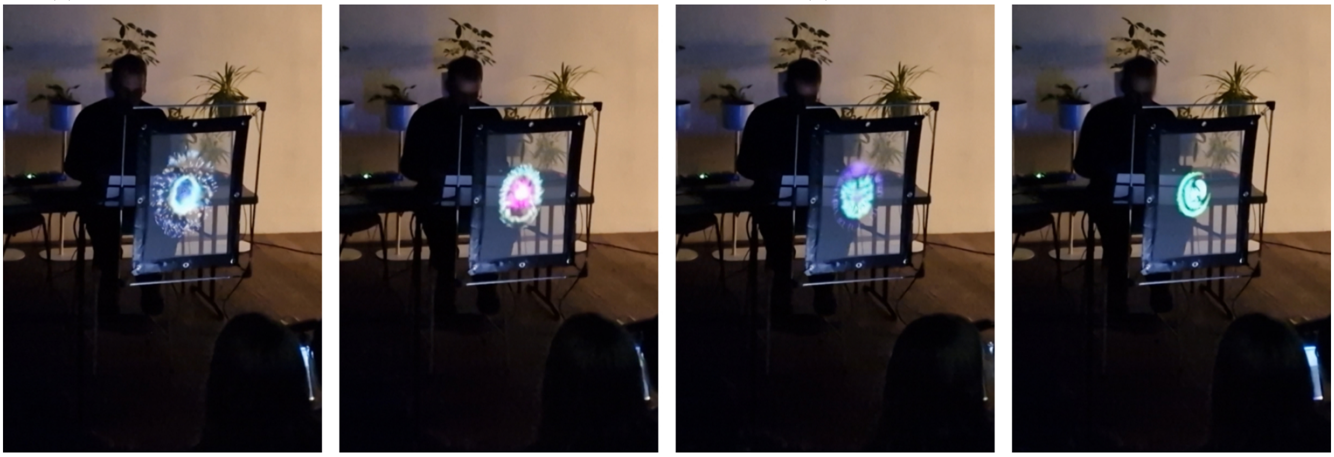

(c) *Pandora's Mycophony*, duo with live coder, August 2023.

Figure 6: Three works created alongside the development of Tölvera spanning May 2023 to present.

proach creates space for the exploration and development of musical and visual ideas, allowing creators/performers to focus on artistic expression rather than dealing with raw data".

5. DESIGN-PRACTICE REFLECTIONS

Although Sections 3 and 4 are sequential in this paper, they occurred concurrently and thus their outcomes and reflections inform each other. Here we reflect first on the workflows that emerged from practice, then we relate these to design reflections.

5.1 Reflections on Artistic Practice

Returning to Table 2 underscores the significance of early collaboration as part of the developmental process, whether via duos (FerroNeural and Pandora's Mycophony) or a trio (MOTHERBIRD). The duos allowed for deeper mapping detail, creating a perception of a more singular system. In contrast, the trio was more distributed in nature, leading to a more peripheral role for Tölvera. Though the duos were in a sense simpler than the trio, compositional complexity increased across species, composition, mapping and projection. In future, more collaborative diversity can be facilitated by making Tölvera programs portable via Taichi's `C/C++` runtime library[15].

We also experimented with different approaches to projection in Tölvera performances. The table setup was more beneficial for the performer than the audience, and the wall setup vice versa. The hologram felt interesting as an alternative approach to ensemble live coding, but needs further simplification to be truly portable. Projection mappings and running headless (without visualisation) will be explored in future, along with physical and tangible interfaces [25].

5.2 Reflections on Design

The version `0.1.0` API design focused on boilerplate code elimination such that pieces could be written quickly with concision, which suited the rapid development of early artistic works. However, initial Tölvera workshops with colleagues have indicated that the terseness of the current API sacrifices discoverability and discourages customisation, and so this will be the main focus of the `0.2.0` design. We expect this will increase the composability of Tölvera programs, for example allowing models to be composed at multiple levels of abstraction, following Levin's notion of multi-scale competency architecture [20]. For example, rather than using the slime mold and flocking models in their complete forms, one could compose together only the sensing capability of the slime mold model and only the movement capability of the flocking model.

As pointed out by one of our reviewers, we have yet to address optimisation, profiling and portability of Tölvera programs, and testing them across multiple platforms and devices (both CPU and GPU). This is recognised as increasingly important work for NIME technologists, in light of the Conference Environmental Statement[16]. Most of the development and practice of Tölvera has taken place on an Apple M1 Max with 64GB RAM, although Tölvera is known to function on Linux and Windows, and on lower capacity devices. Taichi provides mechanisms for profiling[17] and portability[18], which will be explored further after version `0.2.0`. A promising avenue is to export Tölvera programs to single board computers (SBCs) and run them without visualisation, adding Tölvera to the repertoire of tools for high-performance embeddable NIMEs. Further considerations on future developments are considered in the next section.

6. DISCUSSION

6.1 Basal Aesthetics

Basal cognition offers a specific and grounded framework, but we are still required to distinguish enquiry with it into scientific and artistic aspects, where arts are not necessarily oriented in terms of problem-solving or goals. As artists we have been seeking richness and intimacy along tangible interaction principles [17], and familiarity through longitudinal practice. Referring back to Zhang et al's basal sorting algorithm [43], they describe the "unexpected tendencies and capabilities of algorithms", and in Tölvera we propose to frame these as *artistic* tendencies. We can build on their work to draw attention to the *cell's eye view* in Tölvera, the ergodynamics [22] of Tölvera under *perturbation*, and the concept of *algotypes* applied to simple basal processes.

Taking a cell's eye view in Tölvera means that the explicit part of the composition exists in terms of bottom-up, local interactions between particles, that have some aesthetic consequence that only becomes apparent on a global scale. There is thus an implicit level of composition, and composers and performers have the choice of supplementing this with additional omniscient logic and high-level parameters that govern global conditions. From our experience with Tölvera so far we believe that the interplay between bottom-up and top-down decisions produces artistic vitality, if balanced in the right way. Human interventions into Tölvera's basal processes can be likened to experimental perturbations. We can take this idea further by giving users the means to introduce "errors" into particles and models, as Zhang et al did, to observe their impact on ergodynamics from a player's perspective.

6.2 Basal Media for Art & Research

Considering Tölvera as a medium for art and also potentially for research with all kinds of basal materials, its requirements become wide-ranging and diverse. Not least of all it needs to be open-ended like the material it aims to support [39]. We suggest that rather than thinking in mediums, which are ultimately closed, symbolic, and in this case *in silico* domains, we can consider *materiums* [1], as the grey areas where materials and mediums overlap, and consider how to support open-endedness across these borders. From this perspective there are a number of work-in-progress features intended to increase compositionality and the number of modalities available in Tölvera. New `tv.v` models are being developed, for example Particle Lenia [28] and micro neural cellular automata (μNCA) [27].

Lyon et al [21] define a "toolkit of basal cognitive capacities" that includes; orienting response, sensing/perception, discrimination, memory, valence, decision making, behaviour, problem solving, error detection, motivation, learning, anticipation and communication. Some models in `tv.v` already achieves the simpler of these, but the challenge is now to address this wishlist comprehensively. Tölvera's global state dictionary (`tv.s`) will be stress tested to see how well

[15] https://docs.taichi-lang.org/docs/tutorial
[16] https://www.nime.org/environment/
[17] https://docs.taichi-lang.org/docs/profiler
[18] https://docs.taichi-lang.org/docs/taichi_core

it can be used to implement as many of these capabilities as possible. Further, Taichi's NumPy and PyTorch integrations[19] will be explored to give particles and species learning capabilities.

6.3 Diverse Intelligence for NIME

In this work we have focused on intepreting Tölvera from a basal perspective, and it is important to zoom out from there to the broader emerging discipline of diverse intelligence [3, 11, 20, 34], and consider what the implications are for NIME as a field. Recalling Jourdan et al's review of ML in NIME [18], they found that deep learning has "reduced user agency over the models", which are "designed to produce automated processing and the user has rarely access to how these models are formed", further advising that "further study of such trade-offs (including their ethical and political implications) would then be necessary."

This is not to say that ML should be abandoned - as we have seen, IML remains a powerful artistic tool integrated into Tölvera - but rather it should be contextualised within the broader frame of diverse intelligence, to illuminate our collective myopia. With the right tools and context, which we hope Tölvera and this paper provide a basis for, we believe NIME researchers can meaningfully engage with these ideas, and increase the number of collaborations with diverse intelligence researchers.

7. CONCLUSION

In this paper we have introduced Tölvera, a library for composing basal agencies for music interaction, with a number of features that distinguish it from existing artificial life tools. Primarily, it allows for composition of diverse basal embodiments, through its in-built library of interoperable models spanning flocks, slimes, swarms and more, which are easily extendable. Tölvera features globally accessible, declarative style, n-dimensional state and interactive machine learning, making hybridisation and mapping of models trivial. It can also automatically generate OSC endpoints, corresponding Max/MSP and Pure Data patches, and XML and JSON schemas, making integration into musical contexts straightforward. We have also demonstrated live coding, computer vision and sonification capabilities that are in-progress. Through practice-based development, we have only begun to explore Tölvera's musical affordances, and we invite the community to test its limits further. We have also connected Tölvera to the emerging fields of diverse intelligence, basal cognition and agential materiality, to expand the discourse on intelligence within NIME, and promote more-than-human perspectives. Finally, we advocated for more tools like Tölvera in the NIME ecosystem, to support the opening up of new avenues for collaboration with practitioners working across the life sciences.

8. ACKNOWLEDGMENTS

Tölvera has been possible due to collaborations, contributions and support from many individuals and institutions, whom we gratefully acknowledge here.

Paper contributions: Jack Armitage is the creator of Tölvera and wrote and edited this paper, Victor Shepardson is the creator of Anguilla and contributed to the design of Tölvera through many discussions, and Thor Magnusson is the PI of the project and contributed to the writing and editing of this paper. Additional thanks to Victor Shepardson, Isaac Cohen, Robin Morabito and Miguel Crozzoli for feedback on drafts of this paper. We also thank the NIME reviewers, meta reviewer and paper chairs for their valuable and insightful feedback.

Artistic collaborators: thanks to *FerroNeural* collaborator Nicola Privato, *Motherbird* collaborators Jessica Shand and Manuel Cherep, and *Pandora's Mycophony* collaborator Celeste Betancur. In addition, many thanks to Marco Donnarumma for extensive feedback and illuminating discussions.

Intelligent Instruments Lab colleagues: thanks to Halldor Úlfarsson and Sean Patrick O'Brien for holographic display fabrication support, Carla Zimbler for hologauze material sample, Adam Pultz Melbye for ALife conversations and literature exhanges, and Robin Morabito for early user feedback on Tölvera.

Additional collaborators: thanks to *Algorithmic Alphabets* collaborators Luke Iannini (Dynamicland) and Alex McLean (Then Try This) for continuous inspiration and discussion, and to *Sardine* creator Raphaël Forment for Python live coding expertise.

Thanks to the Taichi developers and community for extensive documentation, open source examples and support via Discord.

Workshop and seminar hosts, and participants: thanks to Adam Rokhsar, Sam Tarakajian and community at CuteLab NYC, Bob Sturm and students at MUSAiC KTH Stockholm, Jonas Fritsch and colleagues at AIR Lab ITU Copenhagen, and Martin Kaltenbrunner, Enrique Tomas and students at TamLab Linz.

Residencies: thanks to Elektronmusikstudion (EMS) Stockholm, in particular International Coordinator Gabrielle Karlsén-Beretta, Martin Kaltenbrunner and Enrique Tomas at TamLab Linz, Listaháskóli Íslands' Director of the International Office Björg Stefánsdóttir, and Erasmus+.

The Intelligent Instruments project (INTENT) is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101001848).

9. ETHICAL STANDARDS

The Intelligent Instruments project (INTENT) and has received ethical clearance from the national Science Ethics Committee of Iceland for all research activities, from studies to concerts. No study with human subjects took place as part of this paper, however the project itself has been used in numerous public concerts, installations and workshops. This project and its documentation are open source.

10. REFERENCES

[1] J. Armitage. *Subtlety and Detail in Digital Musical Instrument Design*. Thesis, Queen Mary University of London, Apr. 2022.

[2] J. Armitage, N. Privato, V. Shepardson, S. Skach, A. Temprano, F. Caspe, C. N. Reed, L. Wolstanholme, A. Martelloni, R. Diaz, J. Shier, T. Pelinski, and S. P. O'Brien. Agential Instruments Design Workshop. In *AI Music Creativity 2023*, University of Sussex, Brighton, UK, Sept. 2023. Zenodo.

[3] F. Baluška and M. Levin. On Having No Head: Cognition throughout Biological Systems. *Frontiers in Psychology*, 7, 2016.

[4] C. Betancur and J. Armitage. Pandora's Mycophony. In *Proc. New Interfaces for Musical Expression*, Mexico City, Mexico, 2023.

[19] https://docs.taichi-lang.org/docs/external

[5] C. G. Betancur. AI in live coding environments: Pandora's Dream. In *AI Music Creativity 2023*, University of Sussex, Brighton, UK, Aug. 2023. Zenodo.

[6] D. Bisig, J. C. Schacher, and M. Neukom. Flowspace-A Hybrid Ecosystem. In *NIME*, pages 260–263, 2011.

[7] A. Caillon and P. Esling. RAVE: A variational autoencoder for fast and high-quality neural audio synthesis. *arXiv:2111.05011 [cs, eess]*, Dec. 2021.

[8] B. W.-C. Chan. Lenia - Biology of Artificial Life. *Complex Systems*, 28(3):251–286, Oct. 2019.

[9] M. Cherep and N. Singh. SynthAX: A Fast Modular Synthesizer in JAX. In *Audio Engineering Society Convention 155*. Audio Engineering Society, 2023.

[10] J. Davies and M. Levin. Synthetic morphology with agential materials. *Nature Reviews Bioengineering*, 1(1):46–59, Jan. 2023.

[11] C. Fields and M. Levin. Competency in Navigating Arbitrary Spaces: Intelligence as an Invariant for Analyzing Cognition in Diverse Embodiments. 2022.

[12] R. Forment and J. Armitage. Sardine: A Modular Python Live Coding Environment. In *International Conference on Live Coding*, 2023.

[13] A.-M. Gioti, A. Einbond, and G. Born. Composing the Assemblage: Probing Aesthetic and Technical Dimensions of Artistic Creation with Machine Learning. *Computer Music Journal*, pages 1–43, 2023.

[14] A. Gleadall. FullControl GCode Designer: Open-source software for unconstrained design in additive manufacturing. *Additive Manufacturing*, 46:102109, 2021.

[15] J. Horowitz and J. Heer. Live, Rich, and Composable: Qualities for Programming Beyond Static Text, Mar. 2023.

[16] Y. Hu, T.-M. Li, L. Anderson, J. Ragan-Kelley, and F. Durand. Taichi: A language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics*, 38(6):1–16, Dec. 2019.

[17] R. Jack. *Tangibility and Richness in Digital Musical Instrument Design*. Thesis, Apr. 2019.

[18] T. Jourdan and B. Caramiaux. Machine Learning for Musical Expression: A Systematic Literature Review. In *New Interfaces for Musical Expression*, 2023.

[19] T. Kamatani, Y. Sato, and M. Fujino. Ghost Play - A Violin-Playing Robot using Electromagnetic Linear Actuators. In *NIME 2022*, June 2022.

[20] M. Levin. Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds. *Frontiers in Systems Neuroscience*, 16, 2022.

[21] P. Lyon, F. Keijzer, D. Arendt, and M. Levin. Reframing cognition: Getting down to biological basics. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1820):20190750, Jan. 2021.

[22] T. Magnusson. *Sonic Writing: Technologies of Material, Symbolic, and Signal Inscriptions*. Bloomsbury Academic, 2019.

[23] T. e. a. Magnusson. The Organium: A Library of Technical Elements for Improvisatory Design Thinking. In *Forthcoming*, 2024.

[24] S. Marzban, S. Srivastava, S. Kartika, R. Bravo, R. Safriel, A. Zarski, A. Anderson, C. H. Chung, A. L. Amelio, and J. West. Spatial interactions modulate tumor growth and immune infiltration, Mar. 2024.

[25] A. McLean and L. Iannini. Making Space for Algorithmic Alphabets. In *Undone Computer Science*, 2024.

[26] A. McLean and G. Wiggins. Tidal–pattern language for the live coding of music. In *Proceedings of the 7th Sound and Music Computing Conference*, 2010.

[27] A. Mordvintsev and E. Niklasson. μNCA: Ultra-Compact Neural Cellular Automata, Nov. 2021.

[28] A. Mordvintsev, E. Niklasson, and E. Randazzo. Particle Lenia and the energy-based formulation. https://google-research.github.io/self-organising-systems/particle-lenia/, 2022.

[29] F. Morreale. Where Does the Buck Stop? Ethical and Political Issues with AI in Music Creation. *Transactions of the International Society for Music Information Retrieval*, 4(1):105–113, July 2021.

[30] C. Nordmoen and A. P. McPherson. What the Sensor Knows: More-Than-Human Knowledge Co-Production in Wood Carving. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, DIS '23, pages 779–789, New York, NY, USA, July 2023. Association for Computing Machinery.

[31] K. P. O'Keeffe, H. Hong, and S. H. Strogatz. Oscillators that sync and swarm. *Nature Communications*, 8(1):1504, Nov. 2017.

[32] N. Privato, T. Magnusson, and E. Einar Torf. Magnetic Interactions as a Somatosensory Interface. In *Proc. New Interfaces for Musical Expression*, Mexico City, Mexico, 2023.

[33] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, pages 25–34, 1987.

[34] N. Rouleau and M. Levin. The Multiple Realizability of Sentience in Living Systems and Beyond. *eNeuro*, 10(11), Nov. 2023.

[35] H. Sayama. How Artificial Life Researchers Can Help Address Complex Societal Challenges. In *ALIFE 2021: The 2021 Conference on Artificial Life*. MIT Press, July 2021.

[36] G. Seifert, A. Sealander, S. Marzen, and M. Levin. From reinforcement learning to agency: Frameworks for understanding basal cognition. *Biosystems*, 235:105107, Jan. 2024.

[37] D. Stefani, S. Peroni, and L. Turchet. A Comparison of Deep Learning Inference Engines for Embedded Real-Time Audio Classification. In *Proceedings of the 25th International Conference on Digital Audio Effects (DAFx20in22)*, Vienna, Austria, 2022.

[38] M. Stern and A. Murugan. Learning Without Neurons in Physical Systems. *Annual Review of Condensed Matter Physics*, 14(1):417–441, 2023.

[39] M. Stock and T. E. Gorochowski. Open-endedness in synthetic biology: A route to continual innovation for biological design. *Science Advances*, 10(3):eadi3621, Jan. 2024.

[40] I. van Rooij, O. Guest, F. G. Adolfi, R. de Haan, A. Kolokolova, and P. Rich. Reclaiming AI as a theoretical tool for cognitive science, Aug. 2023.

[41] S. Venkatesh, E. Braund, and E. Miranda. Composing popular music with physarum polycephalum-based memristors. In R. Michon and F. Schroeder, editors, *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 514–519, Birmingham, UK, July 2020. Birmingham City University.

[42] G. Wang and P. R. Cook. On-the-fly Programming:

Using Code as an Expressive Musical Instrument. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 138–143, Hamamatsu, Japan, 2004.

[43] T. Zhang, A. Goldstein, and M. Levin. Classical Sorting Algorithms as a Model of Morphogenesis: Self-sorting arrays reveal unexpected competencies in a minimal model of basal intelligence. Dec. 2023.