A Hyperstitional Machine Appropriating Human Culture in an Evolutionary Fashion

Marinos Koutsomichalis[a]

[a]Department of Multimedia and Graphic Arts, Cyprus University of Technology, 30 Arch. Kyprianou, 3036, Limassol, Cyprus

Abstract

This paper presents an account of an ecosystemic evolutionary pipeline for the generation of original multimedia content without employing fitness functions or other evaluation schemata. It examines an ongoing art/research endeavour that concerns an experimental creative machine to be 'plugged-in' to human culture through the WWW and in order to produce own multimedia content autonomously and unattendedly. The machine employs natural language graphs, as well as intelligent Comprehenders that analyse the retrieved media to further the evolutionary cycle with new queries. It also features a series of algorithmic Composers that mashup and manipulate the retrieved media in various fashions. The overall system is being designed to empirically probe the hypothesis of genuine nonhuman creativity that is built computationaly upon the re-synthesis and the re-appropriation of human culture (through its WWW footprint). The project and the underlying method are announced herein, and a series of technical idiosyncrasies are examined in some detail. Theoretical considerations to the overall approach are further drawn with reference to critical post-humanism.

Keywords

Ecosystemic Evolution, Nonhuman Creativity, Hyperstition, Creative Machine, Multimedia

1. Introduction

Recent literature features numerous resources discussing algorithmic systems for the unattended composition of multimedia content. These range from art/creative endeavors that may or may not involve interaction [1] [2] [3] [4], to bioinformatics [5] and robotics research [6]. The question of unattended and self-generative evolutionary composers has been researched in various contexts such as genre-specific music composition [7], or evolutionary painters [8]. Literature and creative practice is also abundant in algorithmic pipelines synthesising or reappropriating existent (third-party) media content. These range from (historical) examples of music composition employing prepared/found melodies and/or audio snippets [9] to image mashups [10] and multimedia meta-creative systems [11].

Evolutionary creative systems of sort traditionally involve a fitness function or some evaluation schema. They are typically dealt with as meta-heuristics optimization systems, *i.e.*, systems meant to discover those heuristics that are necessary for another subsystem to solve an optimization problem. As discussed in [12], most evolutionary algorithms (EA) for art still follow this approach despite it often being hard, irrelevant, or altogether impossible to define meaningful fitness/evaluation functions in such cases. In genuine artistic contexts, the goal is, more often than not, to generate new and original (or otherwise aesthetically or poetically intriguing) content for the sake of it. In this vein, while 'selection of the fittest' approaches do provide valid and readily available means to implement or evaluate art-related EAs, it remains debatable whether they eventually succeed in generating genuine artistic value in real-life contexts. This is further discussed in [13] and [14], where fitness-based

Joint Proceedings of the ICCC 2020 Workshops (ICCC-WS 2020), September 7-11 2020, Coimbra (PT) / Online
m.koutsomichalis@cut.ac.cy (M. Koutsomichalis)
0000-0002-3876-9064 (M. Koutsomichalis)

EAs are generally shown to concern imitation rather than originally creative behaviour.

While evaluation and selection still govern evolutionary systems study, biological evolution is not exhausted in Darwinian/Lamarckian processes of selection or mutation—see, *e.g.*, [15]—and not all cultural phenomena can be always understood, or described, in terms of meta-heuristics optimisation or problem solving. Insofar as art EAs are concerned, there are some documented cases that eschew or undermine the idea of fitness altogether. Consider, *e.g.*, Biles' jazz melody composer that pivots on an intelligent crossover operator [16], or Dorin's 'interactive' approach that relies on human-driven selection [17]. Another trend is to rather rely on 'endogenous' fitness functions—that is, ones that are defined, and that operate, in some local context rather that with respect to the aesthetic outcome. Even if fitness is still sustained here—both as a concept and as a technical means—a system of sort can no longer be thought of as evolving towards 'fittest'—that is, 'better' in any subjective sense—works of art. A relevant example is Bird's drawing robot where a fitness function rewards local behaviour with respect to pen position [18].

Most importantly, an entirely new paradigm has emerged over the last couple of decades: that of 'ecosystemic' evolution where the focus shifts to the design of an environment, an array of components therein, and carefully designed interactions between the former and latter as well as in-between the components. Components within an ecosystem are typically interconnected so that they can change their environment in some fashion. To give an example, in the Audible Eco-Systemic Interface (AESI) project a network of interdependencies is enacted among individual sound-synthesis subsystems and the external physical space hosting the artistic performance [19]. Accounts of several other art/creative ecosystemic evolution systems can be found in [14] and [20].

All the above mentioned approaches, and in general any algorithmic system for art that is intended as genuinely creative, are still largely thought of with respect to human creativity (even if there is still no consensus on what exactly the latter may stand for, or consist of). The question of a genuinely 'nonhuman' computational creativity—*i.e.*, one that dismisses human notions of creativity altogether—has not been a major research concern hitherto and still lacks integrated treatment. This is, nevertheless, the research focus of this endeavour: to investigate (through design) the hypothesis of a machine that draws upon human culture in order to generate 'nonhuman' art of its own. *Hypersition Bot* is an experimental system that ever-crawls the WWW in order to produce own digital content autonomously and unattendedly. It loosely draws inspiration from the concept of 'hyperstition', brought forth by CCRU's Nick Land and referring to "narratives able to effectuate their own reality through the workings of feedback loops, generating new sociopolitical attractors" (Williams, 2013 as quoted in [21]) or "[…] as ideas [that] function causally to bring about their own reality […] transmuting fictions into truths"[1]. There have been some other attempts to creatively explore this concept and in various fashions—not always artistic, however. Examples are discussed throughout [22].

Hypersition Bot aspires to mashup, transfigure, re-synthesise, remmediate, and re-appropriate—that is, utilise for a different purpose than the intended one—human cultural content with respect to emergent cybernetic orderings and in a hyperstitional fashion. From a technical perspective, the system is an complex multi-modal ecosystemic EA. It does not does employ a fitness function nor any evaluation schemata. It rather comprises a several hardware and software components that intertwine and cross-interact with one another to further the evolution cycle while simultaneously generating multimedia content of various kinds. From an artistic lens, the process is envisioned as speculative (being a hypothesis for how nonhuman creativity could look like), meta-phenomenological (it cannot be reduced merely to phenomenological experiences thereof), and post-geographical (since content

[1]Nick Land in an interview by Delphi Carstens retrieved December 15, 2019 from http://xenopraxis.net/readings/carstens_hyperstition.pdf)

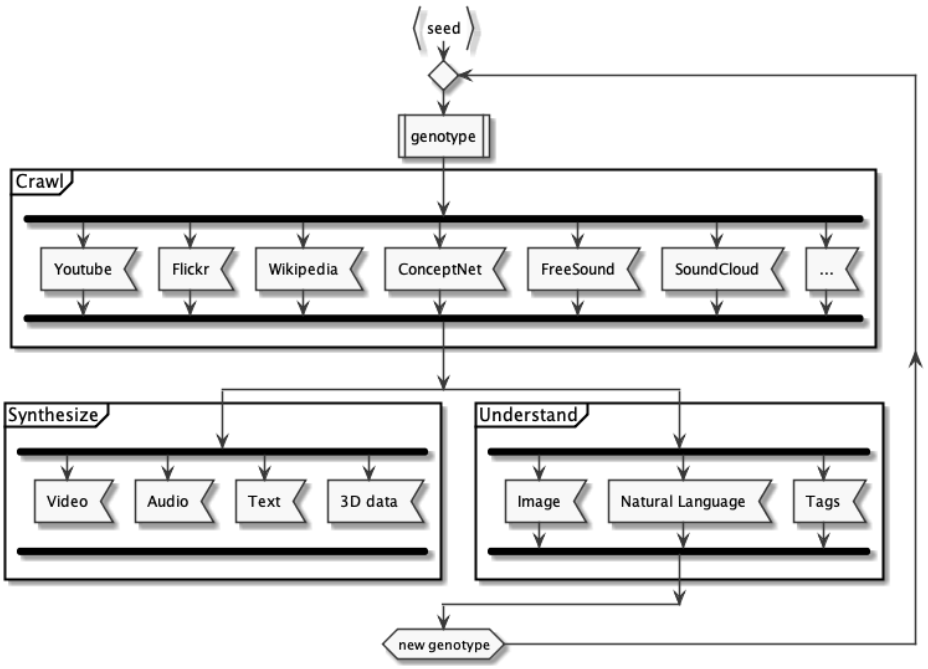

Figure 1: Submodules and their interplay within the evolutionary process.

utilised may be of all possible geographical origins).

Having contextualised the project, the following section outlines the machine in question, overviews the specifics of its implementation, and presents the first incarnation of the work. A Discussion section follows. This treatise sums up with a concluding remarks and notes on future work.

2. Method

The multimedia output of *Hypersition Bot* is the emergent outcome of a complex cybernetic ecosystem that is distributed over several hardware and software modules. Fig. 1 illustrates the main software submodules that have been implemented in a series of programming languages (Python, SuperCollider, Bash). The overall architecture draws on a complex evolutionary database management system that is discussed in great depth in [23]. It additionally features several submodules to perform multimedia synthesis and maintenance. The evolutionary cycle is as follows: a series of Crawlers iterate a genome to retrieve natural language queries and use them to download digital media from User-Generated-Content (UGC) repositories of interest; then, a series of Comprehender submodules analyse the retrieved media to generate a new generation of genotypes, while a series of Composer modules process and mash-up the former to generate multimedia output unattendedly.

In formal terms, the n_{th} generation genotypic population G_n is

$$\mathrm{G_n} := \begin{cases} \bigcup \Phi_\Lambda(P_{n-1}) & : \; n \in \mathbb{N}_{>1} \\ \langle S \rangle & : \; n = 1 \end{cases} \tag{1}$$

where $\langle S \rangle$ indicates the seed—the very first user-defined genome—P_{n-1} the phenotypic population of the previous generation (the digital files retrieved over the WWW in the previous evolution cycle), and

$\Phi_\Lambda : \Lambda^+ \longrightarrow G^+$ a Comprehender submodule mapping phenotypic content of type Λ to new genomes. Comprehenders are of varying complexity/intelligence with respect to the media type Λ they are designed to understand. Three modules of sort are already implemented: (1) Φ_{image} employing the Inception-v3 Deep Convolutional Network [24] and trained on the ImageNet LSVRC-2012 challenge data set [25], (2) Φ_{text}relying on the Rapid Automatic Keyword Extraction (RAKE) algorithm [26] to 'understand' and summarise natural language text, and (3) Φ_{tag} that just converts tags/keywords into genomes.

Individual genomes—as well as populations thereof—are weighted undirected graphs comprising natural language tokens. They have the form $G \equiv (V, W, E)$, where V is a set of vertices $\{v_1, v_2, \dots, v_k\}$ with $k \in \mathbb{N}_{>0}$, $v_n \in U^* \wedge v_k \neq \emptyset$ (U^* denoting all finite (sequences of) words over the unicode character set), W is their scalar weight attributes $\{w_1, w_2, \dots, w_k\}$, $w_k \in [0, 1]$, and E is a (possibly empty) set of pairs $\{v_p, v_\psi\}$ for some $v_p, v_\psi \in V$ and $v_p \neq v_\psi$. All Φ_Λ generate genomic graphs of sort for each one of the files they visit. As explain in detail in [23], at the end of each cycle all available individual genomes are combined to a uniform merger thereof (G_n) so that any cross-associations between their individual edges and weights are resolved. Such a architecture makes it possible to retrieve and manipulate content in many different native languages.

While Φ_Λ are responsible for generating a new genomic population from a given phenotype P_i, a series of Crawler submodules $\lambda_0, \lambda_1, \dots, \lambda_j$

$$\begin{aligned} \lambda_j : \mathrm{G}^+, W &\longrightarrow \bigcup \Lambda^+ \\ Q, R &\mapsto P \end{aligned} \tag{2}$$

are responsible for producing the latter. They do so employing natural language queries Q retrieved over a given genome G to download digital media from the WWW (W), so that the resulting n_{th} generation phenotype P is: $P_n := \bigcup \lambda_j(G_{n-1})$ Note that an individual Crawler λ_j may retrieve digital content of several different types (thus mapping content to $\bigcup \Lambda^+$ rather than Λ^+), so that, *e.g.*, $\lambda_{YouTube}$ retrieves audio, video, and text (user-comments and meta-data). As of writing the system comprises $\lambda_{YouTube}$, λ_{Flickr}, $\lambda_{Shootka}$, $\lambda_{FreeSound}$, $\lambda_{SoundCloud}$, λ_{MLDb}, $\lambda_{Thingiverse}$, $\lambda_{Wikipedia}$, $\lambda_{ConceptNet}$, $\lambda_{WordNet}$ that download audio, video, images, music, prose, lyrics, tags, lemmas, and 3D models from those repositories. It should be noted that Eq. 2 is time-dependent. UGC repositories are volatile so that λ_j will most likely return different results for the same input Q if called at different times.

The above described evolutionary process is implemented in a local network comprising four micro-computers. One of them is responsible for retrieving digital content over WWW, 'comprehending' it, renewing the genome population G for each generation n, and distributing the resulting phenotype P_n among all four. Each of the latter features a series of local helper submodules that handle I/O operation and disk maintenance as needed. Multimedia synthesis is then carried out by a series of Composer modules $T_j : P^+_\Lambda \longrightarrow \Psi$ that manipulate Λ type content to generate new original Ψ type content. As of writing, in all implemented T_j, $\Lambda \equiv \Psi$; there are, however, concrete plans for multi-modal composers. T_j typically process all the available digital files and not just that of the last n_{th} generation. While new content is pushed to the various hardware nodes, older generation P_Λ files are eventually deleted by local maintenance routines. Nevertheless, once the machine has been online for a few evolution cycles, some T_j will almost certainly work on a local pool $L_\Lambda := \bigcup_{i=n-k}^{n} P_{\Lambda i}$ with $k, n \in \mathbb{Z}^+ \wedge k < n$. L_λ comprise content from all $[n - k, n]$ phenotypic populations. The left part of Alg. 1 is an overview of the evolution cycle in pseudo-code.

While the genome mutates in this fashion, a number of Composer submodules mashup or otherwise manipulate the retrieved media files to generate new content. Implementation features a few such submodules, namely: T_{video}, a T_{text}, T_{audio}, and T_{3D}. If Σ is a stochastic operation to select

an element σ_i from L, a simplified model for the video Composer is: $T_{video}(\tau) := \sigma_i(\tau + \rho)$, where $\sigma_i := \Sigma(\bigcup_{n-k}^{n} P_{video}, \tau)$, $\tau \in \mathbb{Z}^+$ denoting discrete time, $\rho \in \mathbb{Z}$ being a random discrete offset (so that video content starts playing back at frame $\tau+\rho$), $\varrho \in \mathbb{Z}^+$ a random discrete time duration after which a new σ_i is to be selected for playback—it should also hold that $\tau+\rho+\varrho \leq \|\sigma_i\|$ (the length of σ_i. The upper right half of Alg. 1 describes this simple mash-up process in pseudo-code. In the actual implementation, Σ is of some complexity, combining chance operations with some hard-coded synthesis rules. The lower right half of Alg. 1 presents T_{3D}—an experimental Composer for solid 3D models drawing on the synthesis pipeline described in [27]. A simplified formal model is $T_{3D} := \bigcup_{j=0}^{n} (Z \circ X \circ R)(O_j)$ with Z, X, R being linear transformations in 3D space that randomly translate, scale, and rotate (respectively) a random selection of individual solid models $\{O_0, O_1, \ldots, O_n\} \subseteq \bigcup_{n-k}^{n} P_{3D}$.

Algorithm 1 Evolution cycle, T_{video}, and T_{3D} in pseudo-code

$G \leftarrow \langle S \rangle$
$\lambda[] \leftarrow \{\lambda_{YouTube}, \lambda_{FreeSound}, \ldots\}$
$\phi[] \leftarrow \phi_{text}, \phi_{image}, \phi_{tag}$
loop
  $P \leftarrow []$
  for $i = 0$ to $\|\lambda\|$ **do**
    $P \xleftarrow{append} \lambda_i(G)$
  end for
  $G' \leftarrow []$
  for $i = 0$ to $\|P\|$ **do**
    $G' \xleftarrow{append} \phi_\Lambda(P[i])$
  end for
  $G \leftarrow \bigcup G'$
end loop

$\Sigma \leftarrow$ a complex stochastic operation
loop
  $\sigma \leftarrow \Sigma(\bigcup_{n-k}^{n} P_{video}, \tau)$
  $\varrho \leftarrow$ a random number in $(0, \|\sigma\|)$
  $\rho \leftarrow$ a random number in $(\varrho, \|\sigma\|)$
  playback σ from ρ to $\rho + \varrho$
end loop

$O \leftarrow$ a random $\subseteq L_{3D}$
$R \leftarrow []$
for $i = 0$ to $\|O\|$ **do**
  $A \leftarrow O[i]$
  random translate A
  random scale A
  random rotate A
  $R \xleftarrow{append} A$
end for
return $\bigcup R$

T_{audio} is an experimental adaptation of a rather complex system for algorithmic mashups that is described in great detail in [28]. The architecture, *inter alia*, comprises a non-real-time machine listening pipeline performing onset detection and spectral feature extraction on all available content. For each audio file $l \in L_{audio}$, it generates a vector $\vec{d}$ registering the particular moments of some notable change (in pitch, rhythm, or timbre), and a feature matrix $\vec{D} \equiv [\vec{c}, \vec{u}, \vec{s}]^T$ with weighted mean frequency $\vec{c} \equiv [c_1, c_2, \ldots, c_k]$, magnitude-weighted variance $\vec{u} \equiv [u_1, u_2, \ldots, u_k]$, and spectral complexity $\vec{s} \equiv [s_1, s_2, \ldots, s_k]$ ($k \in \mathbb{Z}^+$) per some regular time interval. Individual generative 'sonic events' E_k of various different kinds (*e.g.*, (non-)deterministic sequences of shorter sounds, or 'sustained sonic atmospheres') are defined employing audio file fragments and with respect to their associated breakpoints in $\vec{d}_l$ and its origin (the UGC repository they were downloaded from). While E_k are dynamically added to a scheduling queue, an intelligent composition submodule juxtaposes them in real-time and with respect to the feature matrices associated with the audio snippets in use. In this fashion, the particular patterns governing the temporal appearance, the repetition, the duration, and the acoustic localisation are all configured employing features from $\vec{D}$. Some example output (stereo versions) from this experimental Audio Composer can be listened to at https://tinyurl.com/t-audio-examples.

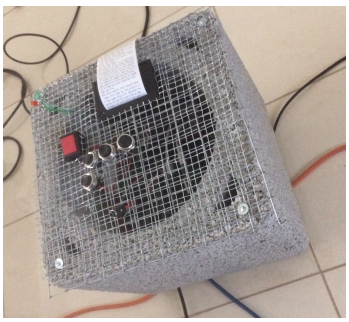

Figure 2: The bot.

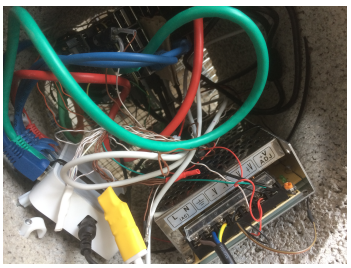

Figure 3: Hardware Submodules.

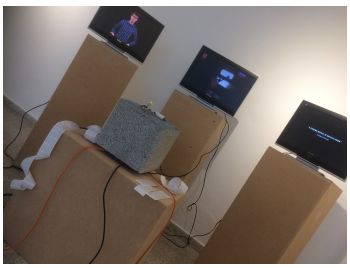

Figure 4: Installation.

T_{text} utilises the 'textgenrnn'[2] system for intelligent character-level text synthesis that is based on a Multiplicative Recursive Neural Network (MRNN) topology. This method is described in great detail in [29]. Given a sequence of input vectors $(\vec{x_1}, \vec{x_2}, \ldots, \vec{x_T})$ a sequence of predictive softmax distribution $P(x_{t+1}|x \leq t)$ is obtained at the output vectors $(\vec{o_1}, \vec{o_2}, \ldots, \vec{o_T})$. The language modelling objective is to maximise the total log probability of the training sequence $\sum_{t=0}^{T-1} logP(x_{t+1}|x \leq t)$. This MRNN topology is ever-trained on every n iteration on some small $l \subseteq L_{text}$. T_{text} is then scheduled to generate new strings of original text at irregular time intervals.

Fig. 2 illustrates the machine in its eventual realisation, with the various hardware submodules hosted in a block of concrete and several cables to interfaces with the WWW, monitor screens, loudspeakers, and other terminals. The machine also features a built-in thermal printer. The overall design is hybrid and rough-hewn, also embodying a certain kind of 'material dialectics'—such an approach towards interface design is further discussed in [30]. Fig. 4 illustrates the machine in its first public showcase in the context of the *Children of Prometheus* international group exhibition that took place in NEME Gallery (Limassol, CY) 2019. In this particular incarnation multimedia synthesis is carried out by just three distinct T_{viceo} and a T_{text} printing out algorithmically generated text every few minutes.

Seen through an ecosystemic lens, *Hypersition Bot* comprises several individual software components of varying complexity that interact with one another and an external environment. This external environment is, in reality, three different overlaid ones: (1) a local network comprising various intertwined hardware and software submodules, (2) the WWW, and (3) the physical space accommodating the *Hypersition Bot* and its generated multimedia content. Accordingly, the proposed architecture is principally grounded on a complex cybernetic network of cross-interactions, inter-dependencies and intertwinments so that its output is emergent, hybrid and distributed—the system's operation cannot be traced, or reduced, to the specifics of its software or hardware modules alone.

3. Discussion

Hypersition Bot has been designed to interrogate and appropriate (*i.e.* using otherwise than intended) the human culture's WWW footprint in a computational and creative fashion. It is envisioned as creative machine that may be plugged-in to a largely human-oriented WWW and creatively re-synthesise media content to bring forth its own alternate, non-human and 'hyperstitional' one. UGC repositories are an excellent way to account both for human culture in its immense trans-geographical and trans-socio-political contingency, as well as for the ways it may be cybernetically 'comprehended' and re-appropriated by machines. YouTube features a few billions of videos of all possible subjects,

[2]Retrieved December 15, 2019 from http://github.com/minimaxir/textgenrnn

themes, and genres uploaded for all possible kinds of purposes by all possible kinds of individuals. Maybe more importantly, it also features meta-data and long threads of structured user comments that further articulate the numerous possible cultural connotations and ramifications of the featured content. Wikipedia is an immense codified and structured database of user-contributed knowledge that covers pretty much all aspects of human existence, from science to popular culture, from history to poetry, and from esoteric religions to design. Soundcloud comprises millions of music works of all possible styles and by all professional, semi-professional, and amateur creators. Music content is further embellished with meta-data and (timeline defined) user comments. UGC of sort unconditionally represent human culture in its sheer eclecticism as well as in the particular ways in which humans themselves interpret, understand, and reflect thereof in all (in)formal, (non)casual and (un)structured, fashions—still, they far-exceed our capacity engage with some significant dimension of them. Without the aid of machines and sophisticated algorithmic techniques it is largely impossible for humans alone to ever make sense of such complex/broad cybernetic phenomena while it is, of course, debatable to what extend we can do so even with the aid of the former.

The creative machine described here is destined to manoeuvre and to appropriate human-oriented cultural content the way it resonates over the WWW, yet in ways humans alone would not be able to pursue. It does so in an ecosystemic fashion, establishing the necessary conditions for emergent cybernetic behaviour to arise. In this vein, *Hypersition Bot* is not meant to imitate human creativity (even if it may accidentally do so at times). It rather celebrates a certain approach towards 'computational poetics'—*i.e.*, an inherently deliberate computational and nonhuman take on multimedia synthesis. But if this so, what would originality and creativity mean in such a context? And how would they relate to human notions thereof? While such affairs cannot be substantially elaborated upon here, there are two important concerns that ought be immediately outlined.

Firstly, the cybernetic method described hereinbefore is primarily meant as an experiment questioning the human authority/exclusivity in both establishing own media culture and in building upon it. As such, *Hypersition Bot* stands together with several other efforts of sorts, that range from like-minded artistic endeavours to the entire philosophical project of critical post-humanism in its various manifestations. Despite their breadth and disparity, the cornerstone of nearly all flavours of critical post-humanism is that humans are ever prosthetic, distributed, and ever produced by, and in relation to, social, environmental, technological, and other nonhuman traits [31, 19–22]. It follows that to further trust/allow machines to re-purpose our own cultural production in computational and nonhuman fashions is actually a very 'human' thing to do—it is what casts us humans in the first place through a post-human lens. The machine's attempt to establish an ouevre of its own can be said to aid to the formulation of a broader and more critical way to understand our species. In the case of *Hypersition Bot* such a stance is deeply echoed to the algorithmic design and its operation that functionally embed the post-human thesis. It is argued that this is only possible through a decentralised ecosystemic design approach that accelerates cross-interactions in between several different 'species', modalities, and (sub)domains. This trait reverberates the herein described machine in all bottom-up and top-down fashions: its made of hybrid and intertwined algorithmic, electronic, and physical components and in a way that brings forth 'material dialectics' of some sort and its operation is emergent and contingent, resonating across digital, analogue, and physical domains and through acoustic, visual, haptic, semantic, and other modalities.

The second concern to delineate relates with the notion of Hypersition. Land's original inspiration traces to Dawkins and his acolytes who popularised the idea that 'memes'—*i.e.* mental elements—control a carrier's thought and behaviour in an ontogenetic fashion and much like genes do to biological bodies [32]. The validity of such a hypothesis is, notwithstanding, questionable in the first place. For instance, Ingold shows that the very claim that some genealogical/inheritance-based mechanism

(exclusively) governs the development of a (biological) organism to some significant extent is feeble: it succumbs altogether under closer examination in that it requisites suitable environmental conditions in the first place [33, 1–17]. Accordingly, he suggests that it is upon the latter we should primarily focus upon when it comes to understand, or to control, how organisms end up being what they are.

It is beyond the scope of this paper to delve in such a debate, of course. Yet, *Hypersition Bot*'s operation is ascribed an additional dimension if seen through a genealogical-contra-ecosystemic prism. The system is specifically designed to fuel a hyperstitional mode of operation: it probes retrieved content to identify and isolate semantic/symbolic links that would link it to other content and so forth, until a hitherto latent (or merely fictional) narrative concretely emerges. The overall system can be said to succeed in such a hyperstitional expedition—at least to some plausible extent—in that it does pursue the cybernetic bearings of our WWW presence and in that it does produce original content re-synthesising them in a generative fashion. However, at the very same time and in tandem with this, it explores the hybrid environmental conditions that cast such a hyperstitional excursion possible: these are, its very own design, WWW and certain UGC topologies within it, the particular cybernetic infrastructures it relies upon to access and to retrieve content of interest, and so on. The machine simultaneously pursues arbitrary congenital bearings waiting to be unfolded, and investigates the conditions that cast such an unfolding possible. It can be then said, that it is made to operate at the crux where ontogenesis and ecosystemic conditioning meet.

4. Conclusion

Hypersition Bot is a creative multimedia system distributed over a complex hybrid network of software/hardware subsystems that creatively explores and re-appropriates the digital footprint of human culture over the WWW in a 'hyperstitional' manner. It autonomously and unattendedly synthesises original multimedia content in a generative fashion and showcases it *in-situ*. 'Original' here stands for appropriated, manipulated and remixed content that attains agency based on how the machine re-purposes it in a cybernetic fashion. The system's architecture is ecosystemic so that its multimedia output is emergent and contingent; it cannot be explained merely in terms of the constituent subsystems. It pivots on technology that is largely designed to 'defy' human creativity in pursue of experimental and nonhuman 'computational poetics'—even if it is yet unclear what it means to be creative in a nonhuman fashion.

The specifics of the various comprising submodules as well as of their interplays and intertwinements are discussed in some detail heretofore. The overall operation is shown to pivot on an ecosystemic evolutionary paradigm that does not employ fitness function or other evaluation schemata of sorts. A couple of concerns surface critical reflection upon the process and the particular poetics at play: (a) the question of challenging human authority/exclusivity in building upon human culture, and (b) the question of 'hyperstitional' behaviour at the crux of ontogenetic and ecosystemic tactics. Affairs of sort are open research questions that call for thorough investigation in both critical and analytical fashions, and, maybe most importantly, speculatively and through the design of relevant creative pipelines. *Hypersition Bot* is such a experimental endeavour, being designed to fumble about the hypothesis of genuine nonhuman creativity.

Considering (a), while it is indeed suggested that *Hypersition Bot* challenges human authority/exclusivity in accessing and building upon human culture in a straightforward manner, and while such a claim is to some certain extent supported pragmatically by means of the machine's overall architecture and multimedia output, this is still a rather bold claim to make and should be taken with a grain of salt. What exactly human authority/exclusivity may stand for in an advanced digital age

is still rather vague—if not ill-formulated. Important investigations in this vein are still an ongoing affair in several subdisciplines such as computational aesthetics, or critical post-humanism. Implementing bidirectional functionality so that *Hypersition Bot* may contribute back content of its own (rather than merely 'consume' human culture) is an important future step towards better formulating the question. So is the design of more complex creative machines of sort. The working hypothesis is that of a machine that (following a long tradition of autonomous algorithmic/generative art) would eventually succeed in transcending human-specific notions of art/creativity altogether, setting out a counter-culture of its own species. In principle, this is the scope of this endeavour: to speculate-through-design on this hypothesis.

Considering (b), it is herein argued that *Hypersition Bot* pivots on the inter-dependency between the exploration of congenital 'genotypic' bearings that wait be unfolded (the ontogenetic dimension), the environmental conditions that cast such an unfolding possible, and the particular ways in which they presuppose, appropriate and establish one another. The overall endeavour could be, therefore, thought of as a structured experiment that both relies upon, and at the same time interrogates, the very conditions that cast 'hyperstitional' creativity possible. Still, what exactly such a creativity is and how it relates with known paradigms of human creativity cannot be answered now—not even properly speculated at. It remains an open research question that needs to be properly formulated and treated in an integrated fashion, both theoretically, and empirically.

5. Future Work

Future research primarily zooms in the implementation of additional Composers, Comprehender, and to a lesser extend, Crawler submodules. T_{audio} and T_{3D} are still in-development and largely experimental. There are concrete plans for a more intelligent T_{3D} pivoting on AI and point-cloud representations of solid geometry that would be inspire by the method described in [34]. There are also plans for cross-modal Composers, *e.g.* $T_{text \rightarrow image}$ or $T_{image \rightarrow 3D}$. Comprehenders submodules Φ_{audio} and Φ_{video} are also needed—even if it is not the latter are rather complex and involved to design. A few additional Crawler modules, *e.g.*, $\lambda_{Instagram}$ or λ_{TikTok}, would also be nice additions. Most importantly, future research zooms in an entirely new class of submodules, that is Uploaders U_j, that would make possible the bidirectional interaction with selected UGC repositories. First priority is for $U_{YouTube}$ and $U_{Twitter}$, with more to follow. When this feature is implemented, *Hypersition Bot* would be granted the right to claim its place in our world, pollinating it with cultural content of its own species.

References

[1] M. Cook, S. Colton, Redesigning computationally creative systems for continuous creation, in: Proceedings of the Ninth International Conference on Computational Creativity (Salamanca, Spain), 2018, pp. 32–39.

[2] A. Das, B. Gambäck, Poetic machine: Computational creativity for automatic poetry generation in bengali., in: Proceedings of The Fifth International Conference on Computational Creativity (Ljubljana, Slovenia), 2014, pp. 230–238.

[3] S. Colton, The painting fool: Stories from building an automated painter, in: Computers and creativity, Springer, Berlin/Heidelberg, Germany, 2012, pp. 3–38. doi:10.1007/978-3-642-31727-9_1.

[4] M. Koutsomichalis, A. Psarra, Computer-aided weaving: From numerical data to generative tex-

tiles, in: Proceedings of the Conference on Electronic Visualisation and the Arts, BCS Learning & Development Ltd., Swindon, UK, 2015, pp. 122–123. doi:10.1145/2641248.2641281.

[5] J. Hartler, M. Trötzmüller, C. Chitraju, F. Spener, H. C. Köfeler, G. G. Thallinger, Lipid Data Analyzer: unattended identification and quantitation of lipids in LC-MS data, Bioinformatics 27 (2010) 572–577. doi:10.1093/bioinformatics/btq699.

[6] M.-J. Kim, T.-H. Song, S.-H. Jin, S.-M. Jung, G.-H. Go, K.-H. Kwon, J.-W. Jeon, Automatically available photographer robot for controlling composition and taking pictures, in: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2010, pp. 6010–6015. doi:10.1109/IROS.2010.5650341.

[7] N. Collins, Musical form and algorithmic composition, Contemporary Music Review 28 (2009) 103–114. doi:10.1080/07494460802664064.

[8] G. Greenfield, Evolutionary methods for ant colony paintings, in: Workshops on Applications of Evolutionary Computation (Lausanne, Switzerland), 2005, pp. 478–487. doi:10.1007/978-3-540-32003-6_48.

[9] M. Koutsomichalis, Catalogue æsthetics: Database in and as music, in: Trends in Music Information Seeking, Behavior, and Retrieval for Creativity, IGI Global, Hersey, PA, 2016, pp. 258–277. doi:10.4018/978-1-5225-0270-8.ch012.

[10] M. Cook, S. Colton, Automated collage generation-with more intent., in: Proceedings of the Second International Conference on Computational Creativity (Mexico City, Mexico), 2011, pp. 1–3.

[11] A. Eigenfeldt, M. Thorogood, J. Bizzocchi, P. Pasquier, Mediascape: Towards a video, music, and sound metacreation, Journal of Science and Technology of the Arts 6 (2014) 61–73. doi:10.7559/citarj.v6i1.129.

[12] C. Johnson, Fitness in evolutionary art and music: what has been used and what could be used?, in: Evolutionary and Biologically Inspired Music, Sound, Art and Design, Springer, Berlin/Heidelberg, Germany, 2012, pp. 129–140. doi:10.1007/978-3-642-29142-5_12.

[13] J. McCormack, Open problems in evolutionary music and art, in: EvoWorkshops 2005: Applications of Evolutionary Computing, Springer, Berlin/Heidelberg, Germany, 2005, pp. 428–436. doi:10.1007/978-3-540-32003-6_43.

[14] O. Bown, J. McCormack, Taming nature: tapping the creative potential of ecosystem models in the arts, Digital Creativity 21 (2010) 215–231. doi:10.1080/14626268.2011.550029.

[15] K. N. Laland, J. Odling-Smee, M. W. Feldman, Niche construction, biological evolution, and cultural change, Behavioral and brain sciences 23 (2000) 131–146. doi:10.1017/S0140525X00002417.

[16] J. A. Biles, Autonomous GenJam: eliminating the fitness bottleneck by eliminating fitness, in: The 2001 GECCO Workshop on Non-Routine Design with Evolutionary Systems, San Francisco, CA, 2001, p. Paper 4.

[17] A. Dorin, Aesthetic fitness and artificial evolution for the selection of imagery from the mythical infinite library, in: European Conference on Artificial Life (Prague, Czech Republic), 2001, pp. 659–668. doi:10.1007/3-540-44811-X_76.

[18] J. Bird, P. Husbands, M. Perris, B. Bigge, P. Brown, Implicit fitness functions for evolving a drawing robot, in: Applications of Evolutionary Computation: EvoWorkshops 2008, Springer, Berlin/Heidelberg, Germany, 2008, pp. 473–478. doi:10.1007/978-3-540-78761-7_50.

[19] A. Di Scipio, 'sound is the interface': from interactive to ecosystemic signal processing, Organised Sound 8 (2003) 269–277.

[20] A. Dorin, A survey of virtual ecosystems in generative electronic art, in: The Art of Artificial Evolution, Springer, Berlin/Heidelberg, Germany, 2008, pp. 289–309. doi:10.1007/

978-3-540-72877-1_14.
[21] S. O'Sullivan, Accelerationism, hyperstition and myth-science, Cyclops 2 (2017) 11–44.
[22] CCRU, CCRU Writings 1997-2003, Urbanomic, Falmouth, U.K., 2017.
[23] M. Koutsomichalis, B. Gambäck, Evolvable media repositories: An evolutionary system to retrieve and ever-renovate related media web content, in: Intelligent Computing-Proceedings of the Computing Conference, Springer, Berlin/Heidelberg, Germany, 2019, pp. 76–92. doi:10.1007/978-3-030-22868-2_6.
[24] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the Inception architecture for computer vision, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (Las Vegas, NV), 2016, pp. 2818–2826.
[25] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, ImageNet: a large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, 2009, pp. 248–255. doi:10.1109/CVPR.2009.5206848.
[26] S. Rose, D. Engel, N. Cramer, W. Cowley, Automatic keyword extraction from individual documents, Text Mining: Applications and Theory 1 (2010) 1–20.
[27] M. Koutsomichalis, B. Gambäck, Generative solid modelling employing natural language understanding and 3d data, in: International Conference on Computational Intelligence in Music, Sound, Art and Design, Springer, Berlin/Heidelberg, Germany, 2018, pp. 95–111. doi:10.1007/978-3-319-77583-8_7.
[28] M. Koutsomichalis, B. Gambäck, Algorithmic audio mashups and synthetic soundscapes employing evolvable media repositories, in: 6th International Workshop on Musical Metacreation (Salamanca, Spain), 2018, pp. 3318–3325.
[29] I. Sutskever, J. Martens, G. E. Hinton, Generating text with recurrent neural networks, in: Proceedings of the 28th International Conference on Machine Learning (ICML-11) (Bellevue, WA), 2011, pp. 1017–1024.
[30] M. Koutsomichalis, Rough-hewn hertzian multimedia instruments, in: International Conference on New Interfaces for Musical Expression (Birmingham, U.K.), 2020-to appear.
[31] P. K. Nayar, Posthumanism, John Wiley & Sons, Hoboken, NJ, 2018.
[32] S. Blackmore, The meme machine, Oxford, Oxford, U.K., 2000.
[33] T. Ingold, Anthropology and/as Education, Routledge, London, U.K., 2017.
[34] S. Ge, A. Dill, E. Kang, C.-L. Li, L. Zhang, M. Zaheer, B. Poczos, Developing creative AI to generate sculptural objects, in: Proccedings of the 25th International Symposium on Electronic Art (Gwangju, Korea), 2019, pp. 225–232.