Generative Solid Modelling Employing Natural Language Understanding and 3D Data

Marinos Koutsomichalis(✉) and Björn Gambäck

Department of Computer Science, Norwegian University of Science and Technology, Trondheim, Norway
marinos@marinoskoutsomichalis.com, gamback@ntnu.no

Abstract. The paper describes an experimental system for generating 3D-printable models inspired by arbitrary textual input. Utilizing a transliteration pipeline, the system pivots on Natural Language Understanding technologies and 3D data available via online repositories to result in a bag of retrieved 3D models that are then concatenated in order to produce original designs. Such artefacts celebrate a post-digital kind of objecthood, as they are concretely physical while, at the same time, incorporate the cybernetic encodings of their own making. Twelve individuals were asked to reflect on some of the 3D-printed, physical artefacts. Their responses suggest that the created artefacts succeed in triggering imagination, and in accelerating moods and narratives of various sorts.

1 Introduction

Computational, computer-aided, procedural, and generative approaches to solid modelling, design, and sculpture are commonly utilized in contemporary research. Willis *et al.* [1] describe *Speaker*, an experimental system converting human voice into three-dimensional contours made of wire. Clune and Lipson [2] have designed an algorithm for evolving 3D objects with a biologically-inspired generative encoding; in their paper they present photos of 3D-printed artefacts and discuss issues related to the physicalization of cybernetic objects. Lehman *et al.* [3] build upon this approach to implement a system that sculpts stylized 3D models. Reed [4] and Horn *et al.* [5] account for evolutionary approaches to vase design. In addition, there is an abundance of resources discussing 'data physicalization' in various contexts [6–12].

Human language has a long history of consistent and significant presence in artistic practice that is not specific to purely language-based media such as literature and poetry—consider e.g. music lyrics, theatrical plays, movie dialogues,

M. Koutsomichalis is supported by an ERCIM "Alain Bensoussan" Fellowship. The authors would like to thank the reviewers for their insightful comments that helped improve the paper.

A. Liapis et al. (Eds.): EvoMUSART 2018, LNCS 10783, pp. 95–111, 2018.
https://doi.org/10.1007/978-3-319-77583-8_7

inscriptions in sculptures, and the various language-driven approaches to conceptual art. Still, examples involving Natural Language Processing are scarce and, maybe not surprisingly due to the technical expertise required, language processing is rarely used in computational solid modelling. Of the few relevant projects that have been documented, it is worth mentioning Stone's *3D printable Alpha Numeric Avatars* [13]—that deterministically store text in 3D-printable objects—and Lee's [14] experiments in 3D-printed ceramic typography.

Digital fabrication and computational solid modelling that a few decades ago were largely confined to engineering and industrial design related disciplines, have become topical to a wide array of functional, scientific and artistic contexts, as the above-mentioned literature demonstrates. The proliferation of open-source hardware/software and amateur oriented 3D-printing has revolutionized practice in areas such as FabLabs [15,16], 'maker' culture [17,18], and DIY (Do It Yourself) personal fabrication [19]. Such research veins have accelerated the DIY, 'hacker' and 'maker' ethos to the general public, resulting in the formulation of various communities that, albeit representing different, and often opposing, ideologies regarding creativity, labour, financial organization, and politics, all pivot on affordable digital fabrication technologies and an ongoing trend calling for the 'datafication' of physical objects [20–23]. Accordingly, and in response to an ever-topical need for 'things' to be fabricated, a series of online repositories has emerged, allowing designers, artists, and digital fabrication enthusiasts to contribute, as well as to retrieve, digital models of all sorts. One of the most populated repositories of the kind, comprising no less than a few millions models, is Thingiverse[1].

This paper introduces an experimental interactive system that generates 3D-printable models such as those illustrated in Fig. 1. It parses and processes arbitrary user-defined textual input, collapsing it to a series of key words or phrases that are subsequently used as queries in order to retrieve 3D-printable models from Thingiverse. The retrieved models are then concatenated into new,

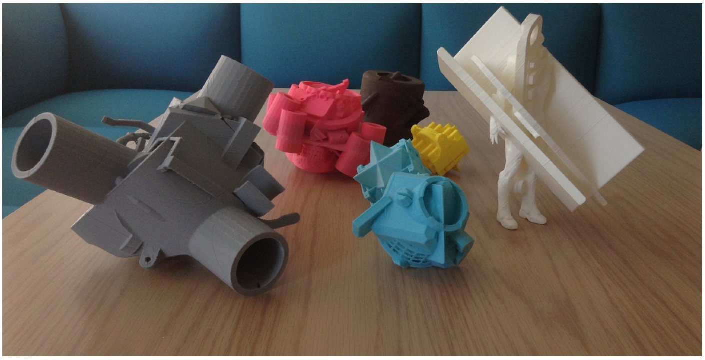

Fig. 1. Examples of computationally generated solid artefacts.

[1] https://www.thingiverse.com.

original ones. The project is intended as both experimental art of a generative kind and, as a hands-on endeavour to probe the materiality of non-numerical kinds of data and of the technologies we rely upon to retrieve, process, and instantiate them. We are not interested in generating aesthetically ordered artefacts, nor in (re-)establishing the original—or some other—specific narrative. Instead, we aim at the production of objects that are original and that, more importantly, look into the affects and effects of their own cybernetic making.

The next section concerns the transliteration of the input text into a bag of 3D models; it gives an overview of the implementation and discusses design decisions and technical shortcomings. Then, Sect. 3 concerns the synthesis and fabrication of original models, and also features photos and illustrations of the resulting artefacts. A theoretical analysis pinpointing aesthetic and ontological affairs comes in Sect. 4. Conclusion and future work then follow.

2 Transliteration Pipeline

The system can be thought of in terms of two consecutive pipelines, as shown in Fig. 2. The first is much more elaborate and concerns the transliteration of the input into a bag of 3D models. The first few stages comprise a module for Natural Language Understanding (NLU). The input text collapses here into a series of 'concepts' to be used as queries to Thingiverse. Each concept comprises correlated words or phrases from the input, and, possibly, synonyms and related terms. The NLU module is implemented in Python, employing the Natural Language Toolkit (NLTK) [24] and relying on either ConceptNet [25] or

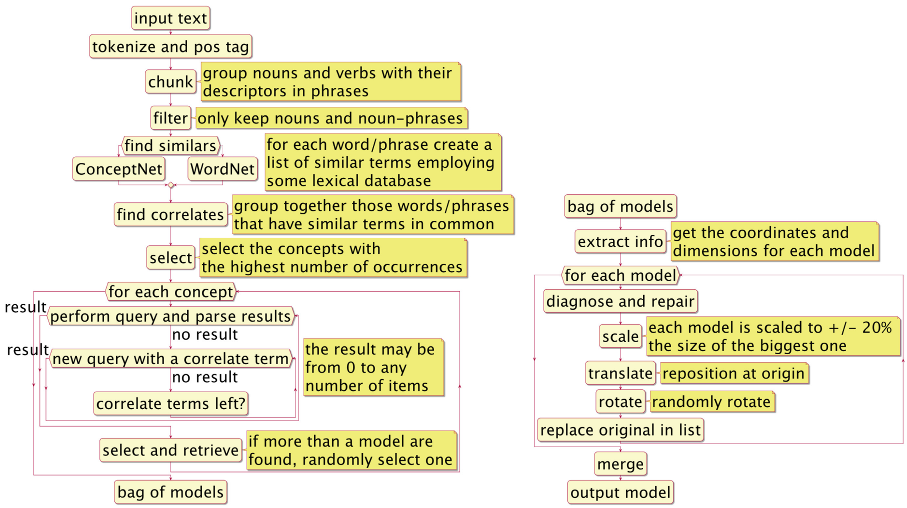

(a) Transliteration pipeline. (b) Synthesis pipeline.

Fig. 2. System architecture.

WordNet [26] for semantic analysis and co-reference resolution. NLTK features direct built-in support for WordNet. Implementation details are given below:

tokenize the input text is segmented and stored as a list of tokens.
part-of-speech tag each token is analyzed and tagged as corresponding to a particular part of speech.
chunk words are chunked together with their descriptors or other auxiliary words to form meaningful phrases.
filter only nouns and simple noun-phrases (containing nouns and adjectives) are kept; common, and therefore nondiscriminatory, English words are also removed using NLTK's built-in 'stopwords' list.
find similars employ either WordNet or ConceptNet to find similar terms for each concept and generate a dictionary where each concept points to an associated list of similars.
find correlates scan the dictionary of concepts, resolve co-references and group together correlates, producing a nested list of concepts, each of which is a list of similars. Note that these lists might contain words that are not present in the original text, but have been retrieved from ConceptNet or WordNet.
select sort with respect to the most often recurring concepts (also taking their similars into account).

The algorithm will keep performing queries at least until either an n (a user defined parameter) number of models have been retrieved, or the list of available concepts is exhausted. For each concept:

query construct the query string and perform a GET request.
parse Thingiverse responds with a JSON-formatted string that has to be parsed: if no (suitable) models have been found, and as long as there are similar terms left, perform a new query using some similar term.
select and retrieve download the model; if more than one model are found, randomly select one.
bag of models append the selected model to a bag of models.

It should be emphasized that it is not at all guaranteed that successful queries return objects that are relevant from a human point of view. As a matter of fact, Thingiverse typically responds with models that are too broadly, or too implicitly, related to the connotations of the original query. This is illustrated in Fig. 3, showing the top results for 'feeling' and 'anxiety' in the category 3D-printing. In the first case, the results look largely haphazard, in the second, too implicit. Such behaviour is a result of Thingiverse's search algorithm, and, more importantly, of the kinds of content contributors choose to upload and of the particular ways in which they name and annotate. Those are largely forged in a chaotic fashion and according to trending popular culture, DIY, fablab, and maker practice, fashion, and all sorts of other cultural veins that shape community-specific interests, as well as everyday life in general. In their turn, ConceptNet and WordNet often respond with concepts that are too broad, too implicit, too context-dependent, and most importantly, occasionally correspond to a non intended reading of the original narrative. For instance, the results for ‘book’, as shown in Fig. 4, include largely arbitrary responses such as ‘caution’ and ‘four of a kind’. Moreover, ‘into’ ranks very high, albeit being a rather generic preposition having no particular significance of its own.

Thingiverse, ConceptNet and WordNet tend to disregard, to a greater or lesser extent, the intended, human-specific, situated and contextualized semantics of a concept and to rather respond according to community-specific cultural trends or the particular idiosyncrasies of the technologies involved. It is exactly because of this that a materialist introspection of the production process becomes possible. The resulting bag of models does not *represent* or *describe* the original text; instead, it exemplifies how the particular algorithms, and databases, involved *interpret* and *understand* it. Yet, even though this system aims to expose, and to creatively experiment with, computational readings and misinterpretations of the sort, totally random results cannot be accepted. Heuristics seem to testify what is required in order to avoid this; in a nutshell: fewer verbs, more nouns, and less abstract queries. The implementation described here primarily draws upon these findings, but also takes into account research demonstrating that it is possible to determine the most important sentences in a text by means of related words or phrases [27,28], to resolve co-reference chains by virtue of semantic databases [29,30], and to successfully convey the meaning of the input text relying solely on nouns and noun-phrases [31,32].

A few examples will illustrate the overall performance of the transliteration pipeline. The program’s log given below is for an excerpt from St. John’s *Revelations* (Rev 9:13–21, The Holy Bible, Revised Standard Version), for $n = 6$ and using ConceptNet. The resulting bag of models is given in Fig. 5a (in all figures, the models are ordered to be read from left to right and from top to bottom).

unto → **found**, **downloaded**
smoke → **found**, **downloaded**
third part → **found**, **downloaded**
god → **found but not usable**, **trying similar**: *angels* → **found**, **downloaded**
silver → **found but not usable**, **trying similar**: *gold* → **found**, **downloaded**
brass → **found**, **downloaded**

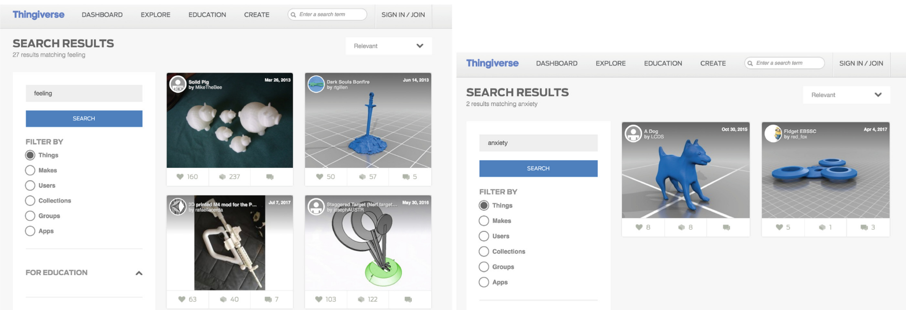

Fig. 3. Thingiverse’s responses tend to be either too broad or too implicit.

Fig. 4. ConceptNet responses can be too broad or relate to unintended readings

(a) Excerpt from *Revelations*: 'angels', 'brass', 'gold', 'smoke', 'third-part', 'unto'.

(b) The *Three Laws of Robotics*: 'existence', 'first', 'human', 'law', 'protection', 'robot'.

Fig. 5. Two example texts transliterated to 3D data.

Apart from the fact that the archaic preposition 'unto' is erroneously identified as noun, the assessed similar terms do indeed correlate, and, in at least two cases, the resulting models correspond to their associated queries: 'smoke'—the model is for a rocket that, of course, can be understood as generating smoke—and 'third part'—the model is obviously part of some broader hybrid.

As another example consider Asimov's famous *Three Laws of Robotics*—taken from "I, Robot" (London, UK: Grafton Books, 1968)—that collapse to 'human', 'first', 'law', 'robot', 'protection', and 'existence', for $n = 6$ and using ConceptNet. In this case all queries are immediately successful, resulting in the bag of models illustrated in Fig. 5b. Here, the extracted terms do summarize the input in a surprisingly evocative manner. However, but for 'robot' and 'human', the remaining concepts are too abstract and, not surprisingly, the resulting models are largely random.

A final example follows. The last two paragraphs of Lovecraft's novel *Shadow over Innsmouth* (New York, NY: Harry N. Abrams, 2016) are used as input, employing ConceptNet and with $n = 19$. The resulting bag of models is shown in Fig. 6 and the program's output log was:

unknown sea → **not found**, **trying similar**: *unheard* → **not found**, **trying similar**: *sea* → **found**, **downloaded**
amidst wonder → **not found**, **trying similar**: *deep* → **found**, **downloaded**
strange → **found**, **downloaded**
sleep → **found**, **downloaded**
father → **found**, **downloaded**
tense → **not found**, **trying similar**: *full change* → **found but not usable**
madhouse → **not found**, **trying similar**: *terror* → **found**, **downloaded**
things → **found but not usable**, **trying similar**: *horror* → **found**, **downloaded**
step → **found**, **downloaded**
ones → **found**, **downloaded**
many → **found but not usable**
abysses → **not found**
innsmouth → **not found**
canton → **not found**
cthulhu → **found**, **downloaded**
splendours → **not found**
stupendous → **not found**
automatic → **found**, **downloaded**
douglas → **found but not usable**
glory → **found**, **downloaded**
lair → **found**, **downloaded**
nthlei → **not found**
reef → **found**, **downloaded**
cousin escape → **not found**
fhtang → **not found**
sanitarium → **not found**
exaltation → **not found**
unto → **found**, **downloaded**
smoke → **found**, **downloaded**
third part → **found**, **downloaded**

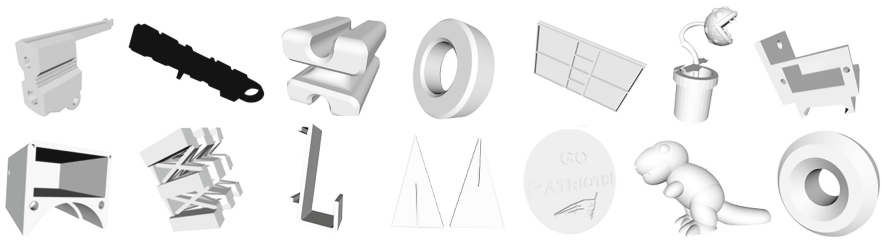

Fig. 6. 3D data for a *Shadow over Innsmouth* excerpt: 'automatic', 'cthulhu', 'deep', 'father', 'glory', 'horror', 'lair', 'ones', 'reef', 'sea', 'sleep', 'step', 'strange', 'terror'.

Not surprisingly, because of the nature of the text, there are many issues here: many queries return no, or not usable, results; nonsensical and too abstract words make it to the list; it is not at all clear how certain words/phrases correlate (e.g. 'unheard'–'sea', 'things'–'horror', 'tense'–'full change', 'amidst wonder' – 'deep'); ConceptNet did not find any similars in far too many cases (e.g. 'abysses', 'canton', 'splendours', 'stupendous', 'cousin escape', 'sanitarium', 'exaltation'); it seems that composite phrases, instead of raising the chances for more relevant results, fail altogether. More, with the striking exception of 'horror', most of the retrieved models are largely irrelevant, or very loosely related, to their generating narratives.

3 Synthesis, Fabrication, Results

Once the input text has been transliterated, the resulting bag of 3D models is passed to the synthesis pipeline which is illustrated in Fig. 2b above. The process is simple and comprises the following steps:

extract info retrieve each model's exact dimensions and position coordinates.
diagnose/repair identify and fix problems (e.g., holes in their mesh).
scale if δ is the largest dimension of each individual 3D model and Δ the largest found δ considering all models, then all models should be proportionally scaled so that δ is always within $\pm 20\%$ of Δ.
translate re-position model to the origin.
rotate each model may, or may not, be randomly rotated a few times around one or more of its axes.
merge concatenate models together.

This part of the system allows for significant improvements. The resulting models are not properly integrated, that is, they are not unique solid objects, but rather aggregates of distinct integrals, so that unnecessary structural elements are still present in areas where they should not be. The output models typically hide complex details, cavities, and micro-masses inside them, which complicate 3D printing. We have experimented with different rotation schemata, so that the input models maintain their original orientation, which is expected to be optimised for fabrication. Such an implementation results in objects that are much easier to fabricate, at the same time, however, their structure is less complicated.

Having explained in some detail the algorithm and its technical offshoots, it remains to discuss the resulting artefacts. Figures 7, 8, 9 and 10 present examples of generated models and their corresponding 3D printed artefacts. In two different occasions and in the context of a semi-formal experimental setting, we presented a total of 12 individuals—both male and female and aged between 25 and 45 years—six of the seven 3D-printed artefacts that are illustrated in Fig. 1 and asked them to write down ideas, thoughts and anything else that would

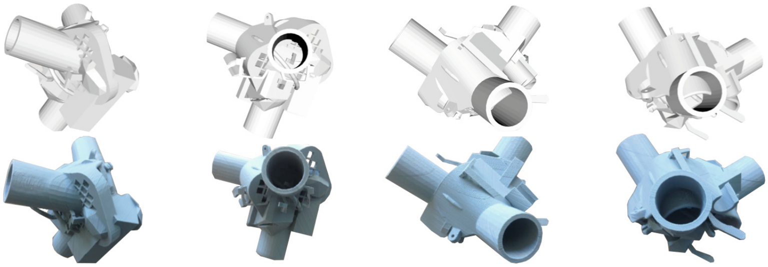

Fig. 7. Input text is an excerpt from Asimov's novel *The End of Eternity* (New York, NY: Tor, 2010, p. 7), referring to the operation of a time-travel machine. Largest dimension: 27 cm. Viewer comments: *"No entrance and no exit. No escape. Something is hiding in there." "A war machine." "Factory. Fly—a magnified head of an insect."*

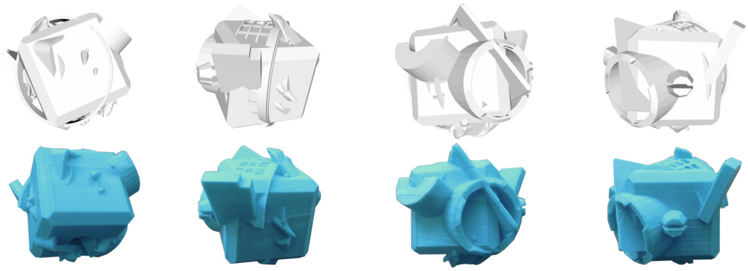

Fig. 8. Input text is an excerpt discussing contemporary representations of finance (Retrieved on January 13, 2018 from www.urbanomic.com/chapter/collapse-viii-nick-srnicek-and-alex-williams-on-cunning-automata/) Largest dimension: 6 cm. Viewer comments: *"Animals that try to escape the prison." "Escape." "A crab mutant created by a scientist." "A house in the forest with the ocean nearby which also has dolphins." "Escape from order towards one thing more amorphous." "It reminds me of a satellite. In general, of something that could be moving in space."* (Translated from Greek by the authors.)

spring to their minds. Some of the responses are rather uninteresting descriptive attempts—e.g. "skull with graph and fragments of antique pillars" regarding the object illustrated in Fig. 9. Several others, however, are imaginative and point at particular moods or even narratives. The most interesting responses are given in each figure's caption.

Those responses prove that, at least up to a certain extent, the artefacts under scrutiny do trigger imagination and do accelerate particular moods, impressions, and even clear-cut narratives. Such responses cannot be attributed neither solely, nor primarily, to subjective idiosyncrasies brought forth by content that is devoid of meaning in a 'Rorschach test' fashion—at least not in principle. Firstly, the above artefacts are typically neither abstract nor devoid of characteristic structure. On the contrary, they often entail recognizable parts of other objects such

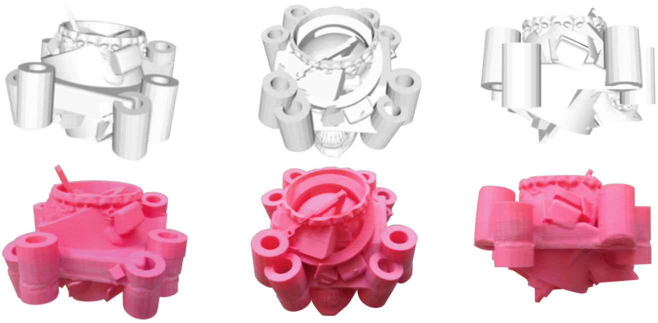

Fig. 9. The input text is a several paragraphs long excerpt from Egan's sci-fi novel *Diaspora* (New York, NY: Skyhorse Publishing Company, 2015, p. 3) describing "*Orphanogenesis*", the birth of a new software-based form of life from a previously unused "*seed*". Largest dimension: 20 cm. Viewer comments: *"Dark and industrial." "A mind full of ideas, but confused. The ideas are heavy for this mind and, moreover, are not that original. They repeat themselves. Obsession?" "Apparently it's a human skull an more precisely a soldier in a field of war." "A man murdered and buried in an underwater facility." "Someone that is trapped in his own universe of discourse which has dismantled him, killed him due to no escape [sic]." "Tetsuo* (Apparently referring to Shinya Tsukamoto's 1989 cyberpunk horror movie "Tetsuo: The Iron Man", the plot of which concerns two men gradually turning to metal.) *(first thought). It's clearly a mixture of organic and industrialized materials. Those tubes look so attached to the skull. It is like magnet pulls them." "Mayas, fear, intense."*

as, e.g. skulls, a human figure, a door sign, a masque, a pair of glasses, or some crustacean's claws, to name a few. Secondly, and more importantly, the above responses cannot be thought of as haphazard since they do often relate to one another and since in some exceptional cases they are even evocative of the original text that generated the object (!).

Indeed, all seven given responses—out of a total of 12—for the object in Fig. 9 suggest a dystopic mood, with some being rather explicit in that we are dealing with some kind of industrial dystopia. To boot, one individual talks of a *"mind"*, *"confusion"* and ideas that *"are not that original"* and *"repeat themselves"*. This is remarkably relevant to the original text, which describes the genesis of a software-based intelligent life form out of a series of existing information patterns and by means of an iterative process. There are no particularly striking correlations to report regarding the object in Fig. 7. Most of the responses for the one in Fig. 8, however, concern either animals, imprisonment, or both. The case of the artefact in Fig. 10 is also remarkable. Apart from the fact that there is a more or less straightforward relevance between all of the given excerpts, some of them are surprisingly relevant to the original text.

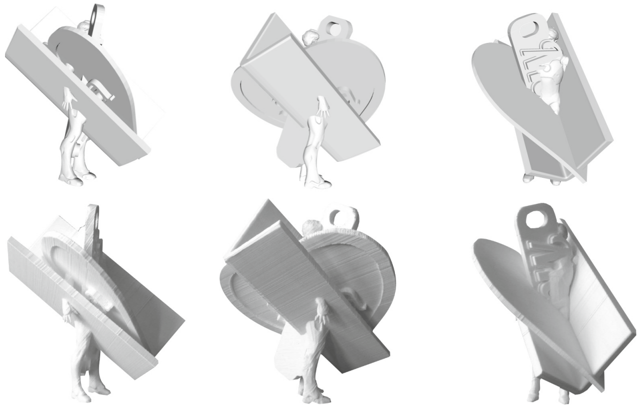

Fig. 10. Input text is the following excerpt from Thoreau's *Walden* (London, UK: Collector's Library, 2010, p. 12): *"The mass of men lead lives of quiet desperation. What is called resignation is confirmed desperation. [...] A stereotyped but unconscious despair is concealed even under what are called the games and amusements of mankind"*. Largest dimension: 25 cm. Viewer comments: *"Drown Valentine."* (Translated from Greek by the authors.) *"A man who's about to go surfing or probably surfman [sic], a super-hero that has just been created by me." "Its about a person that has lost his/her identity. It could also portray a person that tries to bargain the weight of both his heart (emotions) and mind." "Love and humanity." "Film about a man stuck in war (and he's in love?)." "A human's desire for material things." "Heartbroken man, frustration, a door." "A man stricken [sic] by the conjunction of corporate capitalism and love"*

4 Aesthetics, Phenomenology and Post-digital Objecthood

The resulting artefacts seem to satisfy Boden's three criteria for an object to be deemed creative [33]. They are apparently *"new"*, in that, to our best knowledge, similar looking or similarly fabricated objects are not encountered elsewhere. They are definitely *"surprising"*, since they do not look like, nor share similar structural properties with 'ordinary'—both in functional and aesthetic respects—everyday objects. Finally, they are definitely *"valuable"*, at the very least as works of digital art, and, potentially, as crafts worth sharing, gifting or selling. The reactions and the written responses of the individuals who have encountered them clearly testify their potential to trigger human imagination and to evoke particular moods and, exceptionally, narratives.

A closer scrutiny reveals a series of interesting properties as far as both structural and phenomenological aspects of the objects in question are concerned. Albeit physical in the most literal sense of the word, they also reverberate digital concerns and are embedded with the cybernetic encodings of their own making, in that they are both evocative and exemplificatory of 3D-printing technologies and of algorithmic manipulation. In that vein, they celebrate cybernetic 'neomateriality', which according to Paul [34] captures "*an objecthood that incorporates networked digital technologies*" and "*reveals its own coded materiality and the way in which digital processes perceive and shape our world*". The illustrated artefacts are explicitly physical, but they can, nevertheless, be re-produced and re-instantiated *in situ* with respect to localized digital technologies. In addition to being explicitly physical, they are digital information that can be archived, retrieved and re-produced at will, in the very same fashion that digital images [35] or digital audio [36] can. That is to say that they are physical *not in opposition to* their being digital but, instead, *because of, and in addition to*, it. In light of such a hybrid post-digital objecthood, the technicalities discussed in the previous section should be understood as prerequisite conditions bringing forth a materialist introspection at a structural level. That is to say that it is exactly because the transliteration pipeline often results in arbitrary, ambiguous, or micro-culture specific kinds of models that the system eventually produces objects that both 'surprise' and, more importantly, foreground their own cybernetic being.

What particular notions of (post-)digitalness such artefacts are ascribed with still remains to be examined. Philipsen [37] has proposed an analytical schema which is rather adequate here; it understands technologically produced artefacts as hybrids that can be examined from both a digital and a post-digital point of view. He suggests that the digital perspective is largely poetically-concerned and codified with respect to the three aspects of "*cross-disciplinarity*", "*technological essentialism*", and "*artistic creation*", while a post-digital stance takes the ubiquitous presence of digital technology for granted and rather pinpoints readings "*from the outside*", also considering the subjective positions of an audience. From a purely digital stance, we can understand the resulting artefacts as being largely cross-disciplinary: they are simultaneously relevant to, and drawing upon, natural language understanding, information retrieval, computational solid modelling, and digital fabrication technologies. According to Philipsen, a direct offshoot of such a "*refreshingly unorthodox cross-disciplinarity*" is that relevant objects do succeed in foregrounding digital technology in itself.

Philipsen seems to suggest, however, that when dealing with artworks, such technological traits are neither always phenomenologically experienced, nor necessarily readable. Although he does acknowledge *"users"* as a third category specific to interactive art, he argues that there is a certain chasm between *"creators"* (who are aware of technical implications) and *"audiences"* (whose technical understanding stems from para-texts rather than from actual phenomenological encounters with the work). We would argue the validity of this claim, in general and in this particular case. This endeavor both addresses and potentially concerns a broader array of audiences, namely:

(a) ourselves as creators interested in all technical, poetical, and aesthetic aspects of the project;
(b) other artists, scientists, researchers and even hobbyists that may draw upon our method to implement own projects of similar or other sorts;

(c) users who interact with the system in the context of some exhibition setting, in this way generating unique models that may or may not fabricate, or share with others, later;
(d) makers, who do not interact with the system immediately but merely fabricate those already generated models that they appreciate as somehow worth having or producing; and
(e) individuals that visually and haptically engage with the generated objects in some exhibition context.

It may be indeed questionable to what extent technological traits are indeed phenomenologically experienced by audiences of type (e) who may, or may not, be keen on understanding the specifics of the technologies and aesthetic reasoning involved. Audiences (a)–(d), however, are potentially interested in, aware of, and, more importantly, expected to engage, and even tinker with, the various technological and methodical aspects brought forth herein. To boot, the material basis here comprises freely accessible 3D data, 3D-printing, and associated narratives that accelerate openness, participation, and an 'open-source' ethos. The particular kinds of audiences that we indeed address depend, of course, on the particular contexts and exhibition specifics at play. The project as whole, nevertheless, cannot be thought of as *a priori* sustaining an opposition between knowledgeable creators and passive, or technically ignorant, audiences. Here, the know-how is largely open-sourced and audiences are not simply potentially knowledgeable, but also implicitly challenged to *use* the system.

In this way, technological aspects of the work often, if not always, become intrinsic parts of a first-person phenomenological experiences of it. In turn, the social, cultural, and political reverberations of the particular technologies involved are brought forth. Consider that digital fabrication and 3D-printed objects often raise concerns of DIY micro-fabrication, p2p sharing, political activism, 'start-up' entrepreneurship and socially-empowered hacking. Indeed, over the past decade there have been numerous debates and ideological clashes of the various disparate communities that are backing up digital fabrication technologies [38, pp. 25–44]. The artefacts in question can be also thought of as raising concerns of the sort which, therefore, become additional layers of content embedded in their very material fabric.

5 Conclusion

This endeavor has been shown to produce unique, surprising, and potentially valuable artefacts, in a computational fashion, employing arbitrary, user-defined, textual narratives and 3D data retrieved from Thingiverse. The resulting objects celebrate a certain 'neomateriality' in that they articulate cybernetic concerns in concrete physical terms and by means of a hybrid objecthood which is simultaneously digital and postdigital. We understand such artefacts as 'disruptive', both in the sense that they disregard the traditional opposition between the cybernetic disembodiment of information and physical objecthood, and since they are indicative of, and constitutionally relevant to, particular—albeit implicit—(politicized) narratives and contexts. Accordingly, the resulting artifacts have been shown to both foreground the technological specifics of their own making, and to be able of triggering imagination and of accelerating certain moods or impressions. The objects under scrutiny seem to become both evocative and 'disruptive' by virtue of a transliteration pipeline that fails to behave in a definitely transparent and predictable fashion. Indeed, and despite the algorithm's design being largely deterministic, it is never known what lexical associations are forged, what new 'similar' concepts are introduced, what particular kinds of 3D data are retrieved, and, more importantly, how such data are subjectively ascribed new significations (by contributors, users, and commentators).

Using Philipsen's [37] words, the system acquires "*tremendous potential for prompting aesthetic experience*" exactly because of the "*institutional and cultural ambiguity*" it possesses. In this case, creative potential seems to draw on eclectic bags of models that are not easy, or straightforward, to ascribe concrete, or definite, meaning. The transliteration pipeline *must*, then, guarantee that such bags of models always are coherent enough to imply meaning, but, at the same time, ever-eclectic, ambiguous and surprising, so that their structural specifics point back to the very technologies that originally brought forth such ambiguities.

6 Future Work

As discussed, the proposed implementation features a synthesis pipeline which is rather simple and agnostic of context. The input bag of 3D models collapses into a singular solid in a largely haphazard fashion, considering neither the semantic signification of the original queries nor the models' inherent structural properties. An improvement we are considering is that of intelligent geometric transformations that would reflect how the original queries for each 3D model relate to one another, as well as what their signification is in the original input text. This is a non trivial task that involves properly laying out the linguistic relationships at play and somehow mapping them to geometric transformations. We are also considering the possibility of not concatenating 3D models in the first place, but instead, employ deep learning methods so that the algorithm generates original solids that simulate the union, or the intersection, of features of the originals that have to be, of course, first 'learned'. Introducing a module for machine learning could also contribute information to prior stages in the pipeline so that, for instance, the selection of a 3D model among the several returned ones for each query is no longer arbitrary.

References

1. Willis, K.D., Xu, C., Wu, K.J., Levin, G., Gross, M.D.: Interactive fabrication: new interfaces for digital fabrication. In: Proceedings of the 5th International Conference on Tangible, Embedded, and Embodied Interaction, pp. 69–72. ACM, New York (2011). https://doi.org/10.1145/1935701.1935716
2. Clune, J., Lipson, H.: Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In: Proceedings of the 20th European Conference on Artificial Life, pp. 141–148 (2011)
3. Lehman, J., Risi, S., Clune, J.: Creative generation of 3D objects with deep learning and innovation engines. In: Proceedings of the 7th International Conference on Computational Creativity (2016)
4. Reed, K.: Aesthetic measures for evolutionary vase design. In: Machado, P., McDermott, J., Carballal, A. (eds.) EvoMUSART 2013. LNCS, vol. 7834, pp. 59–71. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36955-1_6
5. Horn, B., Smith, G., Masri, R., Stone, J.: Visual information vases: towards a framework for transmedia creative inspiration. In: Proceedings of the International Conference on Computational Creativity, pp. 182–188 (2015)
6. Khot, R.A., Hjorth, L., Mueller, F.F.: Understanding physical activity through 3D printed material artifacts. In: Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems, pp. 3835–3844. ACM, New York (2014). https://doi.org/10.1145/2556288.2557144
7. Stusak, S., Tabard, A., Sauka, F., Khot, R.A., Butz, A.: Activity sculptures: exploring the impact of physical visualizations on running activity. IEEE Trans. Visual. Comput. Graphics **20**(12), 2201–2210 (2014). https://doi.org/10.1109/TVCG.2014.2352953
8. Jafarinaimi, N., Forlizzi, J., Hurst, A., Zimmerman, J.: Breakaway: an ambient display designed to change human behavior. In: Extended Abstracts of the 23rd SIGCHI Conference on Human Factors in Computing Systems, pp. 1945–1948. ACM, New York (2005). https://doi.org/10.1145/1056808.1057063
9. Jansen, Y., Dragicevic, P., Fekete, J.D.: Evaluating the efficiency of physical visualizations. In: Proceedings of the 31st SIGCHI Conference on Human Factors in Computing Systems, pp. 2593–2602. ACM, New York (2013). https://doi.org/10.1145/2470654.2481359
10. Séquin, C.H.: Rapid prototyping: a 3D visualization tool takes on sculpture and mathematical forms. Commun. ACM **48**(6), 66–73 (2005). https://doi.org/10.1145/1064830.1064860
11. Alexander, J., Jansen, Y., Hornbæk, K., Kildal, J., Karnik, A.: Exploring the challenges of making data physical. In: Proceedings of the 33rd SIGCHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 2417–2420. ACM, New York (2015). https://doi.org/10.1145/2702613.2702659
12. Zhao, J., Moere, A.V.: Embodiment in data sculpture: a model of the physical visualization of information. In: Proceedings of the 3rd International Conference on Digital Interactive Media in Entertainment and Arts, pp. 343–350. ACM, New York (2008). https://doi.org/10.1145/1413634.1413696
13. Stone, J.: Systems and methods for transforming and/or generating a tangible physical structure based on user input information. US Patent App. 13/467,713, 23 May 2013

14. Lee, T.: New typographic experience in the post-digital age with 3D printing and ceramics. In: Chung, W., Shin, C.S. (eds.) Proceedings of the 8th International Conference of Applied Human Factors and Ergonomics on Affective and Pleasurable Design, pp. 161–170. Springer, New York (2018). https://doi.org/10.1007/978-3-319-60495-4_17
15. Blikstein, P., Krannich, D.: The makers' movement and FabLabs in education: experiences, technologies, and research. In: Proceedings of the 12th International Conference on Interaction Design and Children, pp. 613–616. ACM, New York (2013). https://doi.org/10.1145/2485760.2485884
16. Posch, I., Ogawa, H., Lindinger, C., Haring, R., Hörtner, H.: Introducing the FabLab as interactive exhibition space. In: Proceedings of the 9th International Conference on Interaction Design and Children, pp. 254–257. ACM, New York (2010). https://doi.org/10.1145/1810543.1810584
17. Tanenbaum, J.G., Williams, A.M., Desjardins, A., Tanenbaum, K.: Democratizing technology: pleasure, utility and expressiveness in DIY and maker practice. In: Proceedings of the 31st SIGCHI Conference on Human Factors in Computing Systems, pp. 2603–2612. ACM, New York (2013). https://doi.org/10.1145/2470654.2481360
18. Kuznetsov, S., Paulos, E.: Rise of the expert amateur: DIY projects, communities, and cultures. In: Proceedings of the 6th Nordic Conference on Human-Computer Interaction, pp. 295–304. ACM, New York (2010). https://doi.org/10.1145/1868914.1868950
19. Gershenfeld, N.: Fab: The Coming Revolution on Your Desktop-From Personal Computers to Personal Fabrication. Basic Books, New York (2008)
20. Dougherty, D.: The maker movement. Innovations **7**(3), 11–14 (2012). https://doi.org/10.1162/INOV_a_00135
21. Toombs, A., Bardzell, S., Bardzell, J.: Becoming makers: hackerspace member habits, values, and identities. J. Peer Prod. **5**, 1–8 (2014)
22. Rayna, T., Striukova, L., Darlington, J.: Open innovation, co-creation and mass customisation: what role for 3D printing platforms? In: Brunoe, T.D., Nielsen, K., Joergensen, K.A., Taps, S.B. (eds.) Proceedings of the 7th World Conference on Mass Customization, Personalization, and Co-creation (MCPC 2014). LNPE, pp. 425–435. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-04271-8_36
23. Groenendyk, M.: Cataloging the 3D web: the availability of educational 3D models on the internet. Libr. Hi Tech **34**(2), 239–258 (2016). https://doi.org/10.1108/LHT-09-2015-0088
24. Bird, S., Klein, E., Loper, E.: Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit. O'Reilly Media, Sebastopol, CA (2009)
25. Speer, R., Havasi, C.: ConceptNet 5: a large semantic network for relational knowledge. In: Gurevych, I., Kim, J. (eds.) The People's Web Meets NLP: Collaboratively Constructed Language Resources, pp. 161–176. Springer, Berlin (2013). https://doi.org/10.1007/978-3-642-35085-6_6
26. Miller, G.A., Beckwith, R., Fellbaum, C., Gross, D., Miller, K.J.: Introduction to WordNet: an on-line lexical database. Int. J. Lexicography **3**(4), 235–244 (1990). https://doi.org/10.1093/ijl/3.4.235
27. Azzam, S., Humphreys, K., Gaizauskas, R.: Using coreference chains for text summarization. In: Proceedings of the Workshop on Coreference and Its Applications, pp. 77–84. Association for Computational Linguistics, Stroudsburg (1999)
28. Barzilay, R., Elhadad, M.: Using lexical chains for text summarization. In: Advances in Automatic Text Summarization, pp. 111–121 (1999)

29. Lee, H., Peirsman, Y., Chang, A., Chambers, N., Surdeanu, M., Jurafsky, D.: Stanford's multi-pass Sieve coreference resolution system at the CoNLL-2011 shared task. In: Proceedings of the 15th Conference on Computational Natural Language Learning, pp. 28–34. Association for Computational Linguistics (2011)
30. Morris, J., Hirst, G.: Lexical cohesion computed by thesaural relations as an indicator of the structure of text. Comput. Linguist. **17**(1), 21–48 (1991)
31. Bergler, S., Witte, R., Khalife, M., Li, Z., Rudzicz, F.: Using knowledge-poor coreference resolution for text summarization. In: Proceedings of 3rd Document Understanding Conference, vol. 3 (2003)
32. Brunn, M., Chali, Y., Pinchak, C.J.: Text summarization using lexical chains. In: Proceedings of 1st Document Understanding Conference (2001)
33. Boden, M.: The Creative Mind: Myths and Mechanisms. Routledge, New York (2004)
34. Paul, C.: From immateriality to neomateriality: art and the conditions of digital materiality. In: Proceedings of the 21st International Symposium on Electronic Art (2015)
35. Groys, B.: From image to image-file–and back: art in the age of digitalization. In: Groys, B. (ed.) Art Power, pp. 83–91. MIT Press, Cambridge (2008)
36. Koutsomichalis, M.: From music to big music: listening in the age of big data. Leonardo Music J. **26**, 24–27 (2016). https://doi.org/10.1162/LMJ_a_00962
37. Philipsen, L.: Who's afraid of the audience? Digital and post-digital perspectives on aesthetics. Peer-Reviewed J. About Post-Digital Res. **3**(1) (2014). http://www.aprja.net/whos-afraid-of-the-audience-digital-and-post-digital-perspectives-on-aesthetics/
38. Hielscher, S., Smith, A.G.: Community-based digital fabrication workshops: a review of the research literature. Technical report, SWPS 2014–08, University of Sussex, Brighton, UK (2014). http://sro.sussex.ac.uk/49214/1/2014-08_SWPS_Hielscher_Smith.pdf. Accessed 13 Jan 2018