

THE TEMPORALITY OF (UN / NON)SELECTION



HANNS HOLGER RUTZ

rutz@iem.at

Institute of Electronic
Music and Acoustics (IEM)
University of Music
and Performing Arts
Graz, Austria

Lisbon
Computation
Communication
Aesthetics
& X

Abstract

One could characterise algorithms by operations of selection—selecting elements, selecting an order between elements, categorising and unambiguously reducing data. It is perhaps through these forms of completion that algorithms exert power, or that some actor attempts to exert power by way of an algorithm. This article proposes that an artistic counter-strategy, a strategy of de-weaponising and aestheticising algorithms, is the conscious exploration of operations of un-or non-selection, that is the interruption of the flow of algorithms, their incompleteness. These operations are elaborated by looking at a number of video pieces, revealing a temporality that cuts across the boundaries of pieces and unpacks the apparent boundaries of algorithms.

Keywords

Algorithms
Violence
Appropriation
Image
Processing

1. INTRODUCTION

With the ethical implications of employing algorithms becoming ever more tangible (cf. Mittelstadt et al. 2016), the field called digital arts can hardly avoid taking standpoints with respect to the increasing exertion of power, the increasing weaponisation of algorithms. Can we simply withdraw to the uninterested position in which algorithms are beyond good and evil, in which the generalised NRA slogan holds: “guns technologies don’t kill people; people kill people”? And if the answer is *no*, can we at the same time resist to surrender the “non-functional” role of art, resist the mounting interpretation—through institutions and funding bodies—of art as superstructure to some assumed political base?

There is, of course, no concluding and uncontested answers, but I propose to look at the temporality of our engagement as human actors with computation processes, and discern specific forms of selection, un-selection, non-selection that could guide us towards a de-weaponisation. The question of responsibility is thus based on the dissolution of the human-machine-opposition (cf. Downey, Dumit and Williams 1995), since we become aligned under the classical statement of the halting problem (cf. Chaitin 1982). There is no general procedure to determine whether an algorithm or a human comes to a halt, the question simply does not make sense. But we can turn it around and ask if we can draw a tableau of breaks and interruptions, not as final selections, as actualisations of some virtual, but as acts of abandonment. Abandonment that could either be understood as un-selection, the movement to a point where something excluded is allowed into the picture, or as non-selection, the non-compliance with the proposition that there is something to be selected *at all*.

As a study object of such operations, I want to look at specific elements of an artistic research project that led up to an exhibition titled *Imperfect Reconstruction* that was realised in 2016 as a collaboration between two sound and digital artists and a stage designer (Rutz 2017). Departing from the perspective that imperfection may well lie at the centre of algorithms and endow them with an intrinsic poiesis, all of the works created during the project in some way or another addressed the question what constitutes imperfection, and how it is possible to make imperfect reconstructions. Semantically, imperfection is not so much understood as a failure of communication, insufficiency or blemish, but as an element of duration and iteration, a resistance to come to a halt.

2. PRELUDE

Before that, it is important to note that by no means these un- and non-selections are specific to digital art. They can be found in any reactive artistic process, that is a process that accepts some amount of empiricism, something that probably holds for the majority of cases. This reactivity is distinctive for site-specific and for installation work, since here by definition the encountered situation configures the artistic intervention. You make a hypothesis about the space, and when you work in-situ, you may discard or modify it. If you attempt to bring a finished piece or a master plan to it, you are inevitably losing the opportunity to create a meaningful interference with the site.

Even further, you may encounter a change in the work *after* it has been installed, and you may choose to accept or reject it. In 2014, I was working on another collaborative installation, *Turbulence*, which featured a space filled with

threads suspended from the ceiling, forming different densities and suggesting specific pathways through the room. What we did not predict is that the threads, made of organic material, would soon form entanglements, even knots, as visitors traversed the space (Fig. 1). At the time, I was talking to a fellow artist about this experience, and how in my eyes it gave an entirely new dimension to the work, which always had been thought with questions of fragility and carefulness in mind. To my surprise, my dialogue partner dismissed this change as accidental to the work and therefore invalid, claiming that the way I incorporated this incident into my discussion of the work was a disingenuous attempt to justify that something unintended and thus unartistic had happened.

From the perspective of this article, the incident was an example of a non-selection: We saw what happened to the installation, but we did not intervene. Or rather, I *did* initially: I visited the gallery several times after the opening and spent hours of disentangling the threads, but I realised soon that I did this not so much to restore a previous order, but as a form of meditation inside the sound installation, a particular way of attending the piece with care. Eventually I decided to only unravel the few sensors that were integrated into the room and needed free movement. I then simply observed the increasing undergrowth and embraced it as something intrinsically anchored in the structure.

Fig. 1

Turbulence. The right side shows a detail of the entanglement resulting from the movement of the visitors.



3. UN-SELECTION

In the case of *Imperfect Reconstruction*, the exhibition space was divided by a three-dimensional mesh structure into an inner and an outer space. One can walk around the outer space which uses the mesh as a contiguous projection surface for a set of connected real-time video works. The mesh is interrupted at two points, allowing one to enter the inner space, characterised by a red surface which shows wandering traces produced by the outer projection shining through the veins of the mesh structure. Eight quadratic screens are installed, half of them mounted inside newly introduced horizontal columns, the other half suspended vertically from the existing vertical columns of the space.

A series of eight-channel video works was developed for this inner space. They try to complement each other as independent miniatures, and each follows a different algorithmic process. Although having been rendered as fixed media, a subtle real-time procedure applies some elements of variation to them, such as shuffling the channels or rotating the image for the horizontally mounted screens. In some series the durations of the eight channels differ, and so in each iteration the relative starting points of the channels are chosen randomly, or a series con-

1 The titles are convenient work titles as the series has never been taken apart into "individuals".


sists of more than eight tracks and in each iteration a subset of eight is chosen. In short, there is never an exact repetition of a situation, and the spatial arrangement encourages one to walk around and see them in different (partial) perspectives and constellations.

The work *Moor*¹ is based on recordings made in a nature reserve of moorlands in January 2016. No special provisions had been made, the footage was collected with a photo camera and without tripod. From a deer stand, one could see all across the moor, and I attempted to make a very slow and steady panoramic movement. It was very cold, and I could not hold my hand still at all times. Every time I noticed my hand was making a too abrupt movement, I stopped and repeated from a slightly earlier position. I anticipated an eventual selection process; I had the vague idea of being able to cut the selected material into one continuous and smooth shot. Everyone who makes sound, video or photographic recordings has this instinct of gathering a surplus, as subsequent software processes are based on operations of selection.

The algorithm applied to this footage entered through a detour—something that, I suspect, is usually what is happening. A month before, I had taken down a show that included a text in white vinyl lettering attached to a wall. Soon I realised that the removal of the letters was tedious and would take a long time, and it would leave the wall, which had been only superficially painted, with white scratches from the underlying colour layer. I interrupted the process, installed a photo camera, and began taking photographs for each successive row—later column—of text removed, turning the wall into an abstract text (Fig. 2). The plan was to create an animated series of the photographs.

From looking at the figure, one can instantly see the change in lighting, due to the fact that the daylight disappeared and the next artist group was already rehearsing in the space with their own lighting. But another problem was more severe: The camera moved slightly between pictures, and for such undertaking the pictures would have to be perfectly aligned. The project with the title *Unlike* (the single word that I left unscratched on the wall) was reflected by a software repository, which I created to undo the camera movement, and the version history allows me to retrace the basic steps of arriving at a particular algorithm.

Fig. 2 1st (top) and 35th (bottom) photograph in the vinyl text removal action.




On 11 January 2016, the software repository was initialised with a number of code pieces taken from another image processing project. One could adjust, ma-

nally, scaling and translation parameters through textual input and see the XOR difference image of two successive photos in the user interface. It was not possible to achieve complete matches. I added a rotational parameter. It did not help, a perspective transform was needed. Two days later I had the perspective transform, and a simple brute force search algorithm was added that would minimise the error. Only it did not converge, when trying to go subsequently from lower to higher resolution. The commit message from 14 January read: "there must be a mistake somewhere, can't believe i can't get any sort of congruence now".

On 16 January, I consulted literature on the subject and found that this was a problem in *image registration*. A group of researchers that were assessing the damage of hurricanes by comparing satellite images were looking for an algorithm to automatically align images that were generally taken from different angles (Thomas, Kareem and Bowyer 2012). They came up with a multi-stage process, and I started implementing the first stage, the *coarse registration*. The idea was to calculate a phase diagram obtained by multiplying the spectrum of the first image with the complex conjugate spectrum of the second image, then go back to the time (space) domain, and the coordinate with the highest pixel value would indicate the translation of the sought transform. The original algorithm would also use a brute force trial of rotation angles before the second image was transformed into the Fourier domain, a step that I did not implement, as rotation seemed irrelevant in my case. With strongly related images, the phase correlation diagram would give one sharp bright spot of only a few pixels extent. Fig. 3 shows such a diagram, with contrast enhanced to show the background structure more clearly that represents all the changes occurring between the two images. The white peak is off-centre towards the top-right, indicating that the camera performed a pan towards the bottom-left between the first and the second picture.

Fig. 3

Phase correlation diagram.
White is positive, black
is negative correlation.



Once these phase diagrams were correctly produced, I un-selected all the sophisticated next steps proposed in the Thomas, Kareem and Bowyer paper, and instead extended the procedure to videos, applying the process pairwise and integrating the translations. It is only now that I came back to the *Moor* piece. The footage being a pan (or actually multiple repeated pans), integrating the translations would result in the image completely leaving the frame to the left in the beginning and to the right in the end of the sequence, so a linear counter motion was added as a measure to keep the image within bounds while still stabilising the motion. I rendered the background black on top of which the translated frames were placed, and something very interesting happened: As the average speed of rotating the camera by hand was not constant, the pan is sometimes "ahead of

time”, sometimes lagging behind. As a result, a new dramaturgy or filmic element is added by a changing vignette. While it is easy to anticipate that this would happen from simply analysing the consequences of the algorithm, the actual effect—the way it unfolds and interacts with the image, the way it shows a particular rhythm—can only be experienced when seeing the resulting video (Fig. 4).

Fig. 4
Still from *Moor* (top) and
assemblage of key frames
(bottom) showing the
relative translations.



Rendering the video required a few more iterations refining the peak-finding function, as the particular material was more sensitive to noise in the phase diagram. But when it was completed, one particular interaction between the material, the context, and the algorithm remained, and it was precisely articulated by the action of un-selecting the subsequent steps of its implementation, un-selecting the full perspective alignment: During the actions of readjusting my arm, the camera was impinged and it produced, for a brief moment, a blurred image and slight rotations around its own axis. The algorithm “works” and “fails” at the same time. It stabilises the translation at the same time as it maintains the perspective distortion which it does not address. The resulting phenomenon transposes the viewed scenery from a credible “immersed” mode of perception—credible in terms of the spatiality of the landscape—into a “mediated” mode of perception, where the landscape becomes almost like a postcard that is being torn apart, or like something separated from the viewer by a lens apparatus which is now revealing its intermediate existence.²

²
The following link leads to a page containing a short video excerpt in which the phenomenon can be witnessed: <https://www.researchcatalogue.net/view/245942/249036> (accessed 02-May-2017).

4. NON-SELECTION

Another piece of the series has the working title *Site*. It also has a past history, albeit a more direct one. Early on in the project, we came up with the German term “Langzeitbelichtung”, or long-term exposure. In this type of exposure, things that happen disparately across time are assembled in a single tableau. For me, it was a metaphor of exposing process, of not targeting a final state that is to be exhibited, but to include all the traces of the processes that can only be understood as ongoing, durational, iterative things into which we “tap” when we frame a project.

Between the beginning of *Imperfect Reconstruction* and its exhibition, I was involved in a different collaborative project, taking place in the public space of a small town. In this project, another artist happened to use an actual long-term exposure process through analogue pinhole cameras. As a partial response to this, I started experimenting with the camera module of the Raspberry Pi platform. I placed a Pi in various places across the town, taking interval photographs and integrating them with an algorithm in a manner somewhat opposed to the analogue integration: Instead of averaging the images over time, I applied a sliding time window median filter that selected or amplified only those pixels that constituted changes in the camera's view. This process produced very curious images that reflected the changes happening over time, changes that are often not obvious to the eye, such as the movement of light, clouds, reflections... (Fig. 5).

Fig.5

Single photo and differential integration of 269 frames.



In *Site*, I was interested in understanding how this process could be translated to moving image or video. Even if one finds a straightforward technical translation, aesthetically this transition is often quite difficult. I started making the first series of images by using the previously developed exposure process, just placing the camera facing one of the gallery's windows, looking to the outside, and leaving it run for a lot longer, recording several thousand images. I then began experimenting with ways of duplicating the sliding window filter as a means to walk through time. The photos being taken every five or so seconds, one starts with a time-lapse video that is quite rapid. I finally applied an audio resampling algorithm, using a band-limited sinc filter, based on time series of each pixel position, slowing down the time-lapse, until it reached a point of sufficient calmness. The particular noisiness and somehow inversion of contrast due to the amplification of differences met another peculiar behaviour: As people walk by the camera's field of view, individual snapshots capture the passersby, while the preceding and successive photos do not show them. There is a reason sinc interpolation is not used in video editing software. It is a resource hungry algorithm, as theoretically the filter kernel is infinite, making it so that every point in time contributes to the interpolated value at any instant. More importantly, the sinc function brings out the Gibbs effect, an over- and undershooting when the input signal sharply changes, as the samples left and right of the slope are alternately weighted with positive and negative coefficients. This phenomenon interacted with the particular recordings of the passersby, producing a strange darkness-brightness oscillation just before their appearance and just after their disappearance in the final video. One gets the impression of contours being "raised" or "falling" cardboard cut-outs. This combines with a particular illuminated green-yellow colour stemming from an unevenness in the camera's RGB gain stages, resembling phosphor (Fig. 6).³

3

The following link leads to a page containing a short video excerpt in which the phenomenon can be witnessed: <https://www.researchcatalogue.net/view/245942/314773> (accessed 02-May-2017).

Example key frames from *Site*, showing the Gibbs oscillation as a person enters the picture, with second and fourth image undershooting to dark, third and fifth overshooting to bright.



It would have been easy to swap the resampling algorithm for another one “more suitable” to video processing, as it would have been easy to adjust the RGB gains or apply a post-production correction. Although none of these elements were planned or prior conceptualised, they gave rise to the particular quality that would be otherwise lost. I simply let go, I let the process run the way “it” came to run, as an act of my own non-selection.

5. TEMPORAL EMULSION

If the previous narration appears to take a long shot at something that may seem rather peripheral, I would like to stress the importance of reading these occurrences as particular breaks cutting into flows of artistic process — with extensive previous histories and successive futures — each of which redirects the flow. It is only due to the limited space that focus was put on *one* instance of un-selection, on *one* instance of non-selection. Further examples of such operations are abundant: The image of the phase diagram was not used (yet). An analogous sound process using the translation estimation was conceived, and a lot of sound material was thus produced, but it was not used (yet). New footage was created with the discovered process in mind, it was not used (yet). In *Moor*, the process for some channels was combined with an imperfect reconstruction of a vector space projection I had seen in Hector Rodriguez’ work *Theorem 8.1* (shown in last year’s *xCoAx*), just implementing the Gram-Schmidt process but then un-selecting the subsequent steps in the algorithm...

Imperfection means not carrying out something to the end. The perspective transform algorithm stems from an analysis of images from before and after hurricane damage, but it is easy to see the immediate military application as well. As artists, we are not only free to enact halting operations, to give up and abandon, perhaps we are also, ethically and aesthetically, obliged to do so. These halting operations are deceptive, because they are not answering to a logical halting problem, and neither do algorithms, as any halting operation gives rise to a bend or leap in the flow that it interacts with.

The undertaking that remains is to draw a more detailed and precise diagram of these operations of un- and non-selection, showing the emulsion of the human time of the artistic process and the machine time of the (imperfect) algorithmic process.

Acknowledgements. Part of this research was undertaken within the project *Algorithms that Matter*, funded by the Austrian Science Fund (FWF), PEEK AR 403-GBL.

REFERENCES

Chaitin, Gregory J. 1982. 'Gödel's Theorem and Information'. *International Journal of Theoretical Physics* 21 (12): 941-954. doi: 10.1007/BF02084159.

Downey, Gary Lee, Joseph Dumit and Sarah Williams. 1995. 'Cyborg Anthropology'. *Cultural Anthropology* 10 (2): 264-269. doi: 10.1525/can.1995.10.2.02a00060.

Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter and Luciano Floridi. 2016. 'The ethics of algorithms: Mapping the debate'. *Big Data & Society* 3 (2). doi: 10.1177/2053951716679679.

Rutz, Hanns Holger, ed. 2017. *Imperfect Reconstruction. An algorithmic project by Lisa Horvath, David Pirrò, Hanns Holger Rutz.* Graz: esc media art laboratory. isbn: 978-3-9503349-1-3.

Thomas, Jim, Ahsan Kareem and Kevin Bowyer. 2012. 'Fast robust perspective transform estimation for automatic image registration in disaster response applications'. In *Geoscience and Remote Sensing Symposium (IGARSS), 2012 IEEE International, 2190-2193.* doi: 10.1109/IGARSS.2012.6351066.