

SIMON CROWE

MICROPOLITICS OF A RECOMMENDER SYSTEM – MACHINE LEARNING AND THE MACHINIC UNCONSCIOUS

INTRODUCTION

In this text I set out to critically examine part of the source code of the recommender system LightFM. To this end I deploy the micropolitics developed by Gilles Deleuze and Félix Guattari, as well as Andrew Goffey's and Maurizio Lazzarato's readings of their micropolitical ideas. I build an argument around Guattari's suggestion that subjectivity is not solely the product of human brains and bodies, and that the technical machines of computation intersperse with what might be thought of as human in the production of subjectivity. Drawing upon contemporary approaches to the nonhuman, machine learning and planetary-scale computation, I develop a framework that situates the recommender system in assemblages of self-ordering matter and links it to historical practices of control through tabulation. While I acknowledge the power of source code in that it always carries the potential for control, in this reading, I impute greater agency to computation. In what follows, rather than reducing it to an algorithm, I attempt to address the recommender system as manifold: a producer of subjectivity, a resident of planet-spanning cloud computing infrastructures, a conveyor of inscrutable semiotics and a site of predictive control.

THE INTERSTICES OF HUMAN AND TECHNICAL MACHINE

In addition to addressing LightFM at the level of its technical workings and infrastructural instantiations, I aim to conceptualise the role these workings play in the ongoing crystallisation of power relations. Deleuze and Guattari offer a useful starting point: everything, for them, is political

and every politics is “simultaneously a *micropolitics* and a *macropolitics*.”¹ They offer the example of “aggregates of the perception or feeling type” and maintain that “their molar organisation, their rigid segmentarity, does not preclude the existence of an entire world of micropercepts, fine segmentations that grasp or experience different things, are distributed and operate differently.”² Much as there are multitudes of molecular variations that escape dominant molar understandings of perception, I argue that there are sinuous historical and machinic paths that criss-cross and circle the linear division of human and technical machine. Further, “there is a double reciprocal dependency between”³ the molar and molecular. Just as multifarious populations are shaped by molar classes, the intertwining histories and relations of humans and technical machines are moulded by the dichotomy’s rigid structure. As a result, some theorisation of humans and technical machines privilege one side of this dichotomy. For instance, in Andy Clark’s notion of the *extended mind*⁴ and Marshall McLuhan’s conception of media as *extensions of man*⁵, media and technical machines are rendered as mere prostheses to human cognition or perception. I favour neither this nor the opposite molar approach of imputing disproportionate agency to technical machines; there is more to explore at the molecular level, in what we might think of as a *machinic unconscious*.

In his book *The Machinic Unconscious*, Guattari posits “a consciousness independent of individuated subjectivity [that] could manifest itself as a component in the assemblages of enunciation, ‘mixing’ social, technical and data processing machines with human subjectivity, but could also manifest itself in purely machinic assemblages, for example in completely automated and computerized systems”⁶. While this partly or fully nonhuman *consciousness* is central to my analysis of LightFM, my reading does not impute consciousness to non-living things. Were it to do so, I might be seen to espouse a form of panpsychism which views “mind [as] a fundamental property of matter itself”⁷; I believe that this position and the questions it raises are beyond the scope of this text. I limit myself to tracing through these “social, technical and data processing machines”,

¹ Gilles Deleuze and Félix Guattari, *A Thousand Plateaus: Capitalism and Schizophrenia*, translated by Brian Massumi, London, Bloomsbury Academic, 2013 [1980], p. 249.

² Ibid.

³ Ibid.

⁴ Cp. Andrew Clark and David Chalmers, “The Extended Mind”, *Analysis*, 58 (1), 1998, pp. 7–19.

⁵ Cp. Marshall McLuhan, *Understanding Media: The Extensions of Man*, Corte Madera, CA, Gingko Press Routledge, 2003 [1964].

⁶ Félix Guattari, *The Machinic Unconscious*, Cambridge, MA, Semiotext(e), 2011 [1979], p. 221.

⁷ Steven Shavero, “Consequences of Panpsychism”, in Richard Grusin (ed.), *The Nonhuman Turn*, Minneapolis, University of Minnesota Press, pp. 19–44, here: p. 20.

semiotic processes that reside in the barren and unmapped hinterlands of human subjectivity. Following Maurizio Lazzarato, I identify “mixed” and fully computational enunciatory assemblages as *proto-subjectivities*. These belong to the register of “non-representational and asignifying”⁸ semiotics: a sign system that operates below the threshold of individuated subjectivity and one that favours pattern over meaning.

Individuated subjectivity, for Lazzarato, is not the sum total of subjectivity; it is produced by the macropolitical process of *social subjection* that assigns subjects “an identity, a sex, a body, a profession, a nationality, and so on.”⁹ Conversely, proto-subjectivity is subjectivity produced by the micropolitical process of *machinic enslavement*. It comprises “a multiplicity of human and nonhuman subjectivities and proto-subjectivities”¹⁰; it arises from the micropolitical process of machinic enslavement, that “dismantles the individuated subject, [...] acting on both the pre-individual and supra-individual levels.”¹¹ Individuated subjectivity is representational, autobiographical and identitarian, giving rise to clearly delineated subject who acts instrumentally upon external objects.¹² Proto-subjectivity is non-representational and pre-individual, capable of emerging in all autopoietic machinic systems.¹³ Critically engaging with proto-subjectivity is by necessity a speculative endeavour calling for non-representational thinking, which Nigel Thrift identifies with the “anti-biographical and pre-individual”¹⁴, with a “vast spillage of things” and with “affect and sensation”¹⁵. LightFM’s source code is an uneasy compromise between the representational, signifying semiotics of individuated subjects and the non-representational asignifying semiotics of technical machines. As such, in this text I must go further than merely explicating the procedures of computation, as I do in the companion text, and address computation in terms of its material instantiations, its histories and its production of affect.

APPROACHING THE MACHINIC UNCONSCIOUS

Andrew Goffey suggested in a recent lecture on the micropolitics of software, that a technological or machinic unconscious might be one that “crosses the histories of programming practices and their shifting relations to the infrastructures that they produce and are produced by –

⁸ Maurizio Lazzarato, *Signs and Machines*, Los Angeles, Semiotext(e), 2014, p. 25.

⁹ Ibid., p. 12.

¹⁰ Ibid., p. 34.

¹¹ Ibid., p. 12.

¹² Cp. Ibid., p. 12.

¹³ Cp. Ibid., p. 80.

¹⁴ Nigel Thrift, *Non-Representational Theory: Space | Politics | Affect*, London, Routledge, 2008, p. 7.

¹⁵ Ibid. p. 10.

[disclosing] the fragmented possibilities of a different relationship to power.”¹⁶ What conceptual tools might help us figure this relationship, to think through these histories and infrastructures, mixed assemblages and proto-subjectivities? We can start by turning our attention, as some contemporary philosophy does, to the nonhuman. One example is object-oriented ontology, a branch of speculative realism that includes thinkers such as Timothy Morton, who would consider a poem, and presumably a piece of code, a nonhuman agent.¹⁷ However, as Jane Bennet notes, object-oriented philosophers refuse the label “materialist”, viewing objects as isolated entities *withdrawn* from other things.¹⁸ Such a position leaves little room for the molecular, micropolitical processes I am concerned with here, many of which are relational or take place close to the undifferentiated level of the machinic phylum. Another more established alternative is Actor-Network Theory, which is primarily concerned with the observation of connections between agents, both human and nonhuman.¹⁹ While an emphasis on nonhuman objects and nonhuman agency is useful – by building on the materialism of Deleuze and Guattari and retaining reference to subjectivity,²⁰ Braidotti’s neo-vitalist materialism provides a better ground for my argument. For her, “all human and non-human entities are nomadic subjects-in-process, in perpetual motion, immanent to the vitality of self-ordering matter.”²¹ We will return to Braidotti later; for now it suffices that autopoietic matter provides a solid ontological substrate for delineating proto-subjectivities.

We now turn to the subject matter of this text, which Adrian Mackenzie addresses in some depth in his recent book *Machine Learners*. Mackenzie’s titular *machine learner* can be both human and nonhuman, or indeed constitute “human-machine relations”²². These human-machine relations are the sites of proto-subjectivities, mixed semiotic assemblages that “inhabit a vectorised space and [whose] operations vectorise data.”²³ I place the vectors and matrices of machine learning in a genealogy of tables as technologies that aid in control. This genealogy stretches from

¹⁶ Andrew Goffey, “Andrew Goffey – Micropolitics of Software”, lecture at Subjectivity, Arts and Data, Department of Media Arts at Royal Holloway, University of London, 2018. Available at: <https://www.youtube.com/watch?v=9bqxsmFo72k> [accessed July 25, 2018].

¹⁷ Cp. Jane Bennet, “Systems and Things”, in Richard Grusin (ed.), *The Nonhuman Turn*, Minneapolis, University of Minnesota Press, pp. 223–240, here: p. 234.

¹⁸ Cp. *Ibid.*, p. 226.

¹⁹ Cp. Bruno Latour, *Reassembling the Social*, Oxford, Oxford University Press, 2008.

²⁰ Cp. Rosi Braidotti, “A Theoretical Framework for the Critical Posthumanities”, *Theory, Culture & Society*, 2018, pp. 1–31. Available at: <http://journals.sagepub.com/doi/full/10.1177/0263276418771486> [accessed March 28, 2019].

²¹ *Ibid.*

²² Adrian Mackenzie, *Machine Learners*, Cambridge, MA, The MIT Press, 2017, p. 6.

²³ *Ibid.*, p. 51.

ancient Mesopotamia²⁴ and encompasses the introduction of tab keys in typewriters²⁵ and the adoption of punch-card tabulating machines²⁶ in turn-of-the-century bureaucracies as well as the relational databases of the 1960s. It is not a straightforward genealogy, however; as Mackenzie notes, for machine learners, the sheer number of dimensions in vector space can “thwart tabular display” and tables can “change rapidly in scale and sometimes in organisation”²⁷. Drawing on Foucault’s account of tables from different eras, Mackenzie argues that the matrices of machine learning mark a return to the Classical or even pre-Classical tables²⁸ that married heterogeneous elements and were structured according to plural and diverse resemblances.²⁹ For example: a matrix of online purchases would place vectors for hair dryers alongside those for garden ornaments; a machine learner, owing to its profoundly flattened ontology, would subject them to the same computation, tracing diverse resemblances through blind repetition.

How might we conceptualise LightFM if it were deployed on a global cloud platform like Amazon? What if, instead of being trained with the ubiquitous MovieLens 100k dataset, LightFM could vectorise the largest ever accretion of user and product metadata on the planet? To aid in answering this question, I will borrow Benjamin Bratton’s model of planetary-scale computation: *the ‘Stack’*. Setting aside the geopolitical intricacies of Bratton’s argument, the stack can be thought of as “a vast software/hardware formation, a proto-megastructure built of crisscrossed oceans, layered concrete and fibre optics, urban metal and fleshy fingers”³⁰. I would argue that planetary-scale machine learning sits at the intersection of the material megastructure of the stack and the asignifying semiotic processes of the machinic unconscious. Bratton’s ‘Stack’ is divided into six layers: ‘Earth’, ‘Cloud’, ‘City’, ‘Address’, ‘Interface’ and ‘User’. These can be placed on a vertical spectrum, rising from molecular to molar – from the geological and chemical, through to the computational all the way up to individuated subjectivity. The micropolitical analysis that follows is primarily concerned with the more molecular ‘Cloud’ layer. However, recalling the double reciprocal dependency between micropolitics and macropolitics, in thinking through the computation that works with vectors and matrices of user

²⁴ Cp. Francis Marchese, quoted in Mackenzie, *Machine Learners*, p. 56.

²⁵ Cp. Susanne Yates, *Control Through Communication*, Baltimore, Johns Hopkins University Press, 1993, p. 80.

²⁶ Cp. James Beniger, *The Control Revolution*, Cambridge, MA, Harvard University Press, 1993, p. 80.

²⁷ Mackenzie, *Machine Learners*, p. 58.

²⁸ Cp. *Ibid.*

²⁹ Cp. *Ibid.*, p. 56.

³⁰ Benjamin Bratton, *The Stack: on Software and Sovereignty*, London, The MIT Press, 2015, p. 52.

meta-data in the ‘Cloud’ layer, we may glean insights into how individuated subjectivity is produced in the ‘User’ layer.

REANIMATING THE CODE

Wendy Chun holds that what we call source code “is more properly an undead resource: a writing that can be reanimated at any time, a writing that haunts our executions.”³¹ I share Chun’s hesitancy to locate agency primarily in source code, which I view as human-readable shorthand with the potential – through multiple translations – to inscribe the palimpsest-like surfaces of computational agency. It is only in that the source code *haunts* its concrete executions that we can read it micropolitically at all. The terse sentences and mathematical formulae Maciej Kula uses to describe LightFM’s algorithm also haunt these executions, but even more spectrally and tenuously than does the code. To illustrate, the following short passage describes the part of LightFM’s algorithm that the source code examples are responsible for implementing: “The latent representation of user u is given by the sum of its features’ latent vectors [...] The model’s prediction for user u and item i is then given by the dot product of user and item representations, adjusted by user and item feature biases”³². Algorithms are divorced from what Goffey calls implementation details: “embodiment in a particular programming language for a particular machine architecture”³³. The Cython source code is as close as my method allows me to get to the micropolitics and asignifying semiotics of computation, but how do I approach it? I heed Mark Marino’s warnings against analysing code aside from its “historical, material, social context” and drawing specious analogies between computation and unrelated social practices or cognitive processes.³⁴ Instead, in what follows, I attempt to speculate beyond the text of the code, to the data structures it references, the materiality of its executions and how these relate to power and control.

Mackenzie observes that many machine learners seek to approximate data by plotting lines and curves through it, or dividing it with lines,

³¹ Wendy Chun, “Wendy Chun – Critical Code Studies”, lecture at the University of Southern California, 2010. Available at: <https://vimeo.com/163282630> [accessed August 25, 2018].

³² Maciej Kula, “Metadata Embeddings for User and Item Cold-start Recommendations”, paper presented in the second workshop on New Trends on Content-Based Recommender Systems co-located with 9th ACM Conference on Recommender Systems, 2015. Available at: <http://ceur-ws.org/Vol-1448/paper4.pdf> [accessed July 27, 2018].

³³ Andrew Goffey, “Algorithm”, in Matthew Fuller (ed.), *Software Studies: A Lexicon*, London, The MIT Press, 2018, pp. 15–20, here: p. 15.

³⁴ Mark Marino, “Critical Code Studies”, *Electronic Book Review*, 2006. Available at: <http://www.electronicbookreview.com/thread/electropoetics/codology> [accessed July 25, 2018].

planes³⁵ and hyperplanes.³⁶ LightFM, however, mainly transforms and compresses a potentially enormous vector space into smaller more easily computable *representations*, on which it bases its predictions. One of the parameters of the *compute_representation* function is the embedding vector for a feature, such as a book genre. The embedding is arrived at based on a matrix, a two-dimensional grid of numbers. In this matrix, each of millions of users is assigned a row and each of hundreds of thousands of books a column; each *cell* where a user and book intersect contains the number “1” if the user has bought the book, otherwise “0”. Now the transformation: from this vast binary matrix an embedding vector of a genre such as ‘software studies’ is produced that points in more or less the same direction as the vectors for other genres of books also bought by software studies enthusiasts. Perhaps the directions of these embedding vectors are among what Mackenzie refers to when he describes the vectorised table as a “machinic process that multiplies and propagates into the world along many diagonal lines.”³⁷ These vectors are arrived at iteratively, through trial and error, in what could be thought of as discretised space and stepwise time. Proto-subjectivities inhabit the discrete space-time of vector computation, just as they reach through a maze of cables, routers and interfaces to the smooth and continuous bodies of users, their *unthinking* habits and gradations of affective intensity.

A non-representational reading of a function called *compute_representation* leaves no space for irony. It leaves little space for users, perhaps more for items, although the function makes no distinction. It expects a sequence of ones and zeros that correspond to embedding vectors, one of which may be the embedding for ‘software studies’. It wants a reference to an existing *representation*, a sequence of floating-point numbers,³⁸ each a *word* 32 bits or binary digits long, arrayed

³⁵ Just as a line is a straight one-dimensional geometric object that extends infinitely in both directions, a plane is a flat two-dimensional object all of whose edges extend infinitely. A point has zero dimensions and can be used to divide a one-dimensional line into two line segments, which can represent classes in the case of a machine learning classifier working with one parameter. A line can be similarly used to divide a two-dimensional parameter space into two classes. For example, if one parameter was human height and the other weight, and the data were plotted on a scatter graph, a straight line could be drawn as a boundary to distinguish the overweight from the non-overweight. The same applies to a plane and a three-dimensional parameter space. As the geometric rules of a plane can be abstracted to arbitrarily high dimensional spaces, a hyperplane of one fewer dimensions than the parameter space can always be used to divide that space. Curved surfaces can be similarly used to classify data at more than three.

³⁶ Mackenzie, *Machine Learners*, p. 212.

³⁷ *Ibid.*, p. 73.

³⁸ An IEEE-754 32-bit floating number comprises three elements: a sign, exponent and fraction. A decimal number can be derived from the floating point representation in the

one after another at a particular numeric address in a memory module in one of thousands of racked servers in a data centre. Execution begins. It marks six 32-bit chunks of memory for later use. It switches context, to the *get_row_start* function in the *features* object; this object is an agglomeration of data and executable instructions sprawling across a bristling microscopic patch of physical memory. It steps through each instruction in this foreign function and remembers the result. It switches back and writes the result to *start_index*, one of the reserved chunks of memory; the same for *stop_index*. One-by one, all bits in all words in the *representation* array are set to 0, switched off. It later, cycle-by-cycle, switches some of these bits on, and sometimes off again, endlessly toggling states; often all bits remain unchanged for an entire cycle as it blindly adds zero to each part of the *representation*.

PLANETARY-SCALE PREDICTION

The notion of deriving a prediction from a representation is firmly situated in the familiar signifying semiotics of individuated subjectivity. And yet *compute_prediction_from_repr* merely calculates the inner product of two vectors, two sequences of numbers, through brute iteration. Mark Andrejevic observes that it is precisely this lack of understanding, reasoning and intuition that gives data-driven prediction its power³⁹ – though *control* may be more apt here. Recall the genealogy of the table alluded to earlier; two-dimensional tables such as timetables are a mark of Foucault’s disciplinary society, which Deleuze suggests has been supplanted by a society of control. I argue that this rupture in the table’s genealogy, the expansion and fragmentation of two-dimensional grids into inscrutable vector spaces, mirrors the fission of enclosed individuals into “‘dividuals’ and masses, samples, data, markets or ‘banks’”⁴⁰. Instead of being placed in a panoptic cell and observed, subjects are divided into row vectors in a panoply of matrices dotted around the ‘Cloud’, and predicted. This control differs from the “purposive influence to a predetermined goal”⁴¹ that James Beniger posits as the seeds of the information society; it is closer to what Luciana Parisi and Steve Goodman call “mnemonic control”: “the power to foreclose an

following way: + or - fraction $\times 2^{\text{exponent}}$. The sign is a single bit that indicates whether the number is positive; the fraction is a 23-bit integer and the exponent an 8-bit integer.

³⁹ Cp. Mark Andrejevic, *Infoglut: how too much information is changing the way we think and know*, New York, Routledge, 2013, p. 21.

⁴⁰ Gilles Deleuze, “Postscript on the Societies of Control”, in *October*, 59, 1992, pp. 3–7, here: p. 5.

⁴¹ Beniger, *The Control Revolution*, p. 36.

uncertain, indeterminate future by producing it in the present”⁴².

How does machinic prediction at the ‘Cloud’ layer entail control at the ‘User’ layer? In constructing vectors that stand in for users, LightFM may be producing categories of subjects, in that these vectors could coalesce or cluster into molar classes. This is a tendency observed by Braidotti, whereby “the neoliberal system finds ways to capitalize also on the marginal and the molecular formations, recomposing them as multiple molarities (i.e. billions of Facebook pages).”⁴³ This can also be figured as purposive movement within the autopoietic matter that makes up the ‘Stack’. Computational proto-subjectivities and their asignifying semiotic chains snake through the ‘Cloud’, ‘City’, ‘Address’ and ‘Interface’ layers, terminating in individuated subjects at the ‘User’ layer.⁴⁴ How might proto-subjectivities apprehend themselves to these subjects? Perhaps as the imperceptible background hum of the ontological power of the future in the present.⁴⁵ If I were to speculate: a subject may feel itself to be acting on blind compulsion. Hunched over a smartphone, through a fog of information fatigue, they may be faintly aware of being nudged toward certain actions, certain products, certain cultural content.

CONCLUSION

In this short text, I have gone some way in posing, if not answering the question of how molecular, machinic processes in a recommender system like LightFM function and relate to power over individuated subjects. I have hinted at how computation may embody a certain kind of agency that feeds into the production of users as subjects. Bratton’s model of *The Stack* aided me in figuring the impingement of proto-subjectivities on molar aggregates such as subjects and users. The idea I have sketched out, of vectorisation and predictive control as a rupture in the genealogy of the table may be an avenue for more extensive research. Although I never primarily attributed agency to the code itself, I could only circumvent the limits of this technical text through speculation. A more empirical approach like Actor-Network Theory may have better elucidated some micropolitical aspects of LightFM, in mapping concrete connections between agents such as computers, programmers and users.

⁴² Luciana Parisi and Steve Goodman, “Mnemonic Control”, in Patricia Clough and Craig Willse (eds.), *Beyond Biopolitics: Essays on the Governance of Life and Death*, Durham NC, Duke University Press, 2011, pp. 163–176, here: p. 167.

⁴³ Braidotti, “A Theoretical Framework for the Critical Posthumanities”, p. 15.

⁴⁴ Cp. Guattari, *The Machinic Unconscious*, p. 51.

⁴⁵ Cp. Mark Hansen, “Our Predictive Condition”, in Richard Grusin (ed.), *The Nonhuman Turn*, Minneapolis, University of Minnesota Press, pp. 101–138, here: p. 125.