Lydia You
GER302—Prof. Rebentisch
9 May 2022
Final Paper

Cult of Engineered (Un)Consciousness: What Are Machines Actually Learning?

When technology gets "smarter", what do humans stand to lose? Technology, and more specifically, computing, promises the ability to process unfathomable amounts of data, which purportedly allows us to make more informed, tailored choices. We are constantly mined for data, and then constrained to use platforms that relentlessly analyze, situate, and reconfigure this data to present us with the most "optimal" choice. It is necessary to explore the implications of the ever-increasing and pervasive application of machine learning to various media, generating a new paradigm of "algorithmic capitalism". With the explosion of volumes of information, the elusive and invisible "algorithm" becomes necessary, and invisible, in every piece of technology we interact with; it is the engine behind every piece of media we are exposed to, but has indelible effects on human agency, socialization, and culture.

At the cusp of the dot-com boom, Negt and Kluge outlined a compelling overview of television as a programming industry that signified a departure from earlier bourgeois ideals of communication, detailing the unidirectional and regulatory function of broadcast media like television and radio. The transition into new forms "privatized" media and the formation of media cartels updated the character of the "consciousness industry". [1] The modern-day algorithmic media platform extends and complicates this definition even further. It is necessary to explore the creation and implications of this new "unconscious" consciousness industry, one

[1] Oskar Negt and Alexander Kluge, *Public Sphere and Experience: Toward an Analysis of the Bourgeois and Proletarian Public Sphere*, vol. 85 (Minneapolis, MN: University of Minnesota Press, 1993), 149.

in which is propagated under the guise of agency and choice made possible through algorithmic prowess, but which erodes humans into homogenous, passive cultural "artefacts"[2].

Algorithm Is King: Peeking Under the Hood

The high-level implementation of a simple recommendation algorithm is described here to illuminate implicit assumptions embedded in the design of a platform like Netflix, as well as assumptions endemic to the field of machine learning. Netflix is a particularly salient example as the company realized relatively early on the potential of aggressive data collection and algorithmic recommendation, pivoting from a mail-order DVD model to a streaming platform in 2007. Moreover, the example of Netflix can be applied broadly to other media platforms such as Spotify, YouTube, and, more disturbingly, Facebook News Feed.

Put bluntly, machine learning is applying basic linear algebra and calculus techniques to real-life datasets with the intent of using existing data to predict future data. It learns, or 'models', real life based on two main presumptions: that real-life phenomena can be sufficiently grasped through many data measurements, and that this data contains patterns that are made visible by applying mathematical axioms that are then propagated out to predict new, unknown results. The recommendation algorithm uses a specific technique of linear algebra called 'Low-Rank Matrix Approximation'. Essentially, the technique 'cross-references' the attributes of movies with attributes of movies that users already like to generate a scalar value that predicts how much that user will like any given movie, even if they haven't seen it before or given it a rating. To do this, it takes two matrices of numbers – one with existing user data (what ratings

[2] Achille Mbembe, in *Out of The Dark Night: Essays on Decolonization* (New York, NY: Columbia University Press, 2021): "The age of the algorithm…turns us all into artefacts and makes redundant a huge part of the muscular power that capitalism relied on."

did users give movies they have watched?) and one with movie data (what is the genre of the movie? Duration? Who stars in it?) – and multiplies them together to obtain a 'product' matrix that will approximate the rating that each user will give to every movie on Netflix. The intuition behind such an algorithm assumes that it can predict a very large set of numbers from a small amount of data; it is precisely trying to approximate the rating that each user would give to each movie, despite only knowing a small number of ratings that users have actually given.[3] This machine learning technique is used widely, from recommendation algorithms to compressing image files.

Several more observations can be discerned from such a design. Flattening all qualities of movies into quantifiable values on a scale is certainly one worth mentioning, especially as a foundational methodology of computing. After all, digital machines can only speak in numbers, yet give a machine enough numbers and it can fill in the picture in ways not even visible to humans. We assume we can quantify each movie: on a scale of 1-5, how romantic is it? How funny? How exciting? How heart-warming is it? And then we go further: how likely would you be to recommend this movie to a friend undergoing a breakup? How "socially acceptable" are the characters? And on and on. Every quality, and every scenario, is quantified and capitalized on; the data that can be generated is infinite. You then get a list of movies that have been maximized for enjoyment since the algorithm splits hairs; perhaps one movie is recommended over another because the overall enjoyment score is higher by a hundredth of a point (what do these points of "enjoyment" even signify?). These hairbreadth differences become reified over time, as the user never gets to see what is *not* presented. What, then, does it mean to have a

[3] This description is based on my own implementation of Low Rank Matrix Approximation in a Princeton Computer Science class, but the technique has been extensively researched and documented. This is a good reference from a Stanford introductory ML class: https://web.stanford.edu/class/cs168/l/l9.pdf.

number determining how much you will enjoy a movie? Even more so, one could argue that using a point system to quantify "enjoyment" is rather arbitrary and one-dimensional, not to mention that the enjoyment is based upon minute calculations premised on the fact that users like movies for their specific definable attributes (romanticness, funniness) rather than for other more holistic qualities. The algorithm also has the capability of cross-referencing users it has identified as having similar taste to recommend movies that you will probably like, as evidenced by the many recommendation features on Amazon, Goodreads, Spotify, etc. that will curate recommendations continually as you add more ratings or buy more items (i.e. other users who liked this book also liked these other books). This begs the question: whose choice is it anyways?

The Syntax of 'Choice'

It seems as if every oddly specific category exists on Netflix, like "*Feel-good Foreign Comedies for Hopeless Romantics"* or "*Scary Cult Mad-Scientist Movies from the 1970s".* There is something uncannily human-like, but decidedly mechanically accurate and deterministic, in the way the algorithm recommends and categorizes films. To understand why, it is important to delve into the 'grammar' and 'vocabulary' of how movies are described internally. Madrigal describes how, armed with a document cheekily dubbed "Netflix Quantum Theory", Vice President of Product Todd Yellin devised a plan to hire people to watch movies and assign 'tags' to them. Tags range from traditional descriptors of genre, happy or sad endings, lead actors, etc. to more ambiguous and uncommon attributes, like plot conclusiveness or how visually striking the cinematography is. The tags are then assigned scalar values from 1 to 5. From these tags, the software teams at Netflix devise descriptive 'microgenres', i.e. "Mind-bending Cult Horror Movies from the 1980s", that would form the basis of the recommendation system. In other

words, this process allows Netflix to pinpoint exactly what kind of movie – even if you couldn't put it into words yourself – that you like.[4]

If we were ever under the illusion that movies were defined by something larger and undefinable, like their narratives or a sense of artistic vision, that illusion has been thoroughly punctured with the proliferation and reification of these 'microgenres', of which there were exactly 76,897 of them when Madrigal's article came out nearly a decade ago. The concept of tagging movies is interesting in that it digests the content of films retroactively into a set of agreed-upon terms. Notably, tags can be stacked, creating endless combinations of long, quirky descriptors. Each tag is considered equally important, meaning that the more tags, the more meaningless each individual tag becomes. Thus, movies become dismembered and metabolized into word-salad that succeed only in capturing the effect of films.

The Making of the Unconscious Consciousness Industry

Negt and Kluge introduce the term "consciousness industry" to refer to the whole of the complex phenomena of public-sphere mass media. The consciousness industry is distinguished from precedents like Adorno and Horkheimer's "culture industry" and Haug's "illusion industry" as it defines human consciousness as the primary "raw material and mode of appropriation"[5]. The programming of mass media has changed dramatically since the era of public television and radio. Mass communication then was intended for a large heterogeneous audience whose specific likes and needs were largely unknown to the media companies. Negt and Kluge emphasize the

[4] Alexis C. Madrigal, "How Netflix Reverse-Engineered Hollywood," *The Atlantic* (Atlantic Media Company, December 22, 2018), https://www.theatlantic.com/technology/archive/2014/01/how-netflix-reverse-engineered-hollywood/282679/.
[5] Negt and Kluge, "The Context of Living as the Media Cartel's Object of Production", *Public Sphere and Experience,* 149.

unidirectionality and mass anonymity of the audience as defining factors in this medium, describing an asymmetric relationship. Because the audience was unknown, anonymous, and massive, the programming often relied on generic and abstract content that was regulatory in function, rather than communication that inspired engagement or response. They considered television an industrial enterprise, conjuring the image of media companies assembling and manufacturing a fleet of sleek, uniform cars on an assembly line, but rather than cars, it is the immediate and complete image of a cultural product that is projected on a screen. [6]

The shift to "privately organized" media (including cassettes, videodiscs, and cable broadcasts) brought about a new character to programming. The new media were able to transcend genericness and speak directly to each individual, and thus shifted the target of the object of production: "their programs do not merely comprise an abstract all-purpose package ('to whom it may concern') but are able to make individualized needs, the needs of target groups, and thereby whole contexts of living, the object of a focused opportunity for exploitation"[7]. For the first time, human consciousness—our preferences, attention spans, and needs—were the object upon which the industry calculated upon, as evidenced by the focus on "demand" as an impetus for production. The unidirectional relationship remained, however, as audiences were still unable to respond.

The introduction of algorithms into our media platforms has complicated this asymmetrical relationship further, and it is useful to consider Negt and Kluge's points to define a new form of "unconscious consciousness industry". Users perceive that platforms like Netflix and Spotify open up the possibility of a bi-directional communication channel, where each person can communicate their specific likes and dislikes and receive a curated, individualized

[6] Negt and Kluge, "Public-Service Television", 100-103.
[7] Negt and Kluge, "The Context of Living as the Media Cartel's Object of Production", 155.

experience rather than generic broadcast media. The promise of infinite choice (as well as the increasingly fine-grained specificity of the choices, as with the long microgenre names), coupled with the illusion of agency in selecting personalized choices, obscure the unconscious mechanisms of data collection and comparison happening behind the scenes, the details of which were roughly sketched out above. The algorithm collects not just data one would expect, like ratings or genre preferences; it collects "secondary" numbers on how long you spend watching each film, at what times of day, in what locales, and what you watch before and after. It is an ever-present companion in our daily rituals of eating, sleeping, and entertainment. The algorithm engages more senses than we knew we had; a subconscious, visceral, and even biological form of extraction for the purpose of the commercial.[8] It can be said to transcend traditional divides of private and public communication; rather, a recommendation algorithm engages in a mixed private-public model that privately mines data on an individual basis and situates it in a public network of millions of other users.

If cassettes and videodiscs brought our media from the big screen into our homes, then the introduction of streaming platforms and personal devices has brought it even further, onto our persons and even into our minds. Negt and Kluge describe the perceived "immediacy" and "completeness" of television, denoting how films seem finished and "life-like" on the screen, obscuring the many workers and hours that went into production. The algorithm, then, encourages the uncanny feeling that it knows you better than you know yourself, being able to predict what you will like and hiding other options, increasing the sense of these "ready-made" entertainment options. The ubiquity of these apps on personal devices that are getting

[8] Negt and Kluge describe how deception and obfuscation of the cultural and societal effects the consciousness industry has is a key propagating force of it: "In contrast, the late-capitalist enterprise that runs the consciousness industry must conceal the social effects it has, and must strictly underestimate its influence so as not to erode the base of its private initiative" (159).

increasingly more portable and proximal to our bodies, from big screens to family televisions to smartphones, contributes to the perception of easily accessible and infinite media. Users often forget, however, the recommendations are optimized on only a few key axes – time spent on the platform and ratings given – promising a rather empty and passive experience of 'consuming' media. In this way, algorithmic platforms promise more and better, but unconsciously detract from users' autonomy and "sensuality" by secretly mining data, processing it in ways hidden from the public, and re-packaging it back to each user, discouraging exploration and concealing alternatives.

This new sensuality is further qualified: "What Lothar Hack has described as the 'new immediacy' has emerged: needs must be satisfied immediately, because one perceives that the material possibility for satisfying them is present."[9] The directive to satisfy cravings of entertainment has driven the devolvement of classical media and art into new, bite-size packets of "content". Content primarily seeks to fulfill a purpose and is solely produced according to the "demands" of its audience. Content is delivered directly to its buyer. Media platforms nowadays engage primarily in content over all other forms of media. Even Netflix executives refer to their process of dismembering each movie into thousands of tags as the "quanta", or "little packets of energy"[10], that comprise a movie, thoroughly degrading it from its whole form into tiny packets of content.

Algorithmic culture, which may have seemed an oxymoron decades ago, has become so ingrained that the cultural products it pushes at its users become metabolized into real lifestyle effects; it takes on a social program. The algorithm does not just collect and reconfigure our data; it goes beyond "passive" machinery and compels action. Nowhere is this more evident than in

[9] Negt and Kluge, "The Context of Living as the Media Cartel's Object of Production", 156.
[10] Madrigal, "How Netflix Reverse-Engineered Hollywood".

the formation of "aesthetics" on image-and-video-sharing platforms like Instagram and TikTok. "Aesthetics", as used in its slang pop-culture context, refers to the formation of online communities around certain "lifestyles" and appearances that are defined and disseminated via visual media. One popular example is "cottagecore", which primarily consists of a "romanticized interpretation of western agricultural life" with principles of simple living, harmony with nature, and self-sufficiency. Its visual signifiers largely consist of images of girls donning flower crowns lounging in a green meadow, a loaf of homemade sourdough resting on a tablecloth, motifs of strawberries and mushrooms, foraged plants, vintage books, and the like.[11] The primary method of enrolling new "members" is through exposure to more and more images of the aesthetic. Montages with flicks of similar motifs scroll by, all with the imperative that they are "aesthetic"; that is, if one partakes in the aesthetic and captures it, the appearance must be beautiful. One can display that they are a member and identify fellow members of the aesthetic through images and appearances alone. Aesthetics compel its members to take action and build a lifestyle around its images: "People are united as individuals, but they experience this union through capital. They can recognize each other only via this apparatus. Collectivities are formed, but without self-regulated interpersonal relationship; forms of satisfaction develop, albeit passive ones. That which exists is organized, but there is no autonomous activity"[12]. There is a powerful force of homogenization here. The algorithm seeks to diagnose each user with their own "aesthetic" based on collected data, in which each individual is fed a steady diet of images and encouraged to realize those images into real life, conforming to the archetype of that "aesthetic".

[11] A detailed explanation of "cottagecore", its founding philosophies, key attributes, history, etc. can be found here: https://aesthetics.fandom.com/wiki/Cottagecore. The wiki contains thousands of pages describing various aesthetics in detail.
[12] Negt and Kluge, "The Context of Living as the Media Cartel's Object of Production", 158.

It may be fitting to update Kluge and Negt's definition of media cartels – that is, media corporations engaged in consolidating profits to themselves through various means[13] – to "data cartels" in the age of the algorithm. Netflix has a staggering amount of data on media consumption, yet this data is siloed from other companies and the algorithm is proprietary and its details kept secret. It is not surprising, then, that Netflix has decided to invest heavily in original productions in recent years, with nearly 40% of its total annual spending allocated to making original works. With unmatched stores of data, Netflix is positioned well to give viewers exactly what they want and tighten their stranglehold on movie production and dissemination. Modern day media platforms have combined the movie studio, production company, marketing agency, and movie theater all into one.

We have reached a point in culture where software triumphs hardware. No longer is the statement that "the most important investment interests are those of the electronics industry which produces the hardware"[14] true; rather, access to immense stores of data and a proprietary, aggressive recommendation algorithm define the modern media cartel.

Culture Industry Reconsidered, Reconsidered

The phrase "if you are not paying for it, you become the product" has become so trite and misunderstood that it is important to break down the modern "culture industry". Adorno and Horkheimer define a "culture industry" that is distinct from a culture that would naturally arise out of the masses, one in which "conformity has replaced consciousness" and produces standardized, passive people and middling products. Their culture industry describes humans as

[13] Negt and Kluge, "The Context of Living as the Media Cartel's Object of Production", 150.
[14] Ibid., 151.

debased. Despite the transparent deception of the culture industry as a profit-driven structure, participants willingly turn a blind eye for fleeting gratification.[15]

We are now engaged in a self-reinforcing cycle too, as made apparent by the example of algorithm-driven media platforms. Humans are compelled to participate, lest they miss out on accessing media altogether; their participation in the 'data economy' becomes reinforced as companies learn more about them to deliver a better "service" to them. Humans will change their behaviors and inputs to cooperate with machines, as evidenced by the odd and decidedly inhuman scramble for search-engine optimization and willing subjection to the whims of the all-mighty algorithm; machines collect these inputs, often without humans' knowledge or consent, and feed it into an algorithm; the algorithm discretizes and analyzes this input and reproduces it back to us, diluted with other people's input and your own historical data. We act on this output, even share to others what we're watching and listening to via the social features of these platforms. The companies behind the algorithms then create new cultural products based upon our inputs to feed back to its customers. It is not so much that algorithms reproduce social norms – rather, algorithms appropriate, digest, and then recycle existing norms, but also generate a new kind of conformity. They homogenize, make us passive, while giving the illusion of unprecedented agency and choice ("individuality itself serves to reinforce ideology"[16]). Humans – our consciousness, our every decision, even our biological whims– are capitalized upon and become artefacts, rendered ineffective by our machines.

[15] Theodor W. Adorno and Max Horkheimer, "The Culture Industry: Enlightenment as Mass Deception," in *Dialectic of Enlightenment* (New York/London: Continuum, 1976), pp. 120-167 and Theodor W. Adorno, "Culture Industry Reconsidered," in *The Culture Industry: Selected Essays on Mass Culture* (London: Routledge, 1991), pp. 98-106.
[16] Adorno, "Culture Industry Reconsidered", 101.