

Durham Research Online

Deposited in DRO:

25 May 2022

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Amoore, Louise (2022) 'Machine learning political orders.', *Review of International Studies* . pp. 1-17.

Further information on publisher's website:

<https://doi.org/10.1017/S0260210522000031>

Publisher's copyright statement:

Open Access. This article is distributed under the terms of the Creative Commons Attribution License (CC-BY 4.0), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

SPECIAL ISSUE ARTICLE

Machine learning political orders

Louise Amoore* 

Department of Geography, Durham University, Durham, United Kingdom

*Corresponding author. Email: Louise.amoore@durham.ac.uk

(Received 6 May 2021; revised 22 October 2021; accepted 17 December 2021)

Abstract

A significant set of epistemic and political transformations are taking place as states and societies begin to understand themselves and their problems through the paradigm of deep neural network algorithms. A machine learning political order does not merely change the political technologies of governance, but is itself a reordering of politics, of what the political can be. When algorithmic systems reduce the pluridimensionality of politics to the output of a model, they simultaneously foreclose the potential for other political claims to be made and alternative political projects to be built. More than this foreclosure, a machine learning political order actively profits and learns from the fracturing of communities and the destabilising of democratic rights. The transformation from rules-based algorithms to deep learning models has paralleled the undoing of rules-based social and international orders – from the use of machine learning in the campaigns of the UK EU referendum, to the trialling of algorithmic immigration and welfare systems, and the use of deep learning in the COVID-19 pandemic – with political problems becoming reconfigured as machine learning problems. Machine learning political orders decouple their attributes, features and clusters from underlying social values, no longer tethered to notions of good governance or a good society, but searching instead for the optimal function of abstract representations of data.

Keywords: Politics; Machine Learning; Algorithm; International Order; Rules

Introduction: ‘We need to learn the features’

During the political campaigns for the US presidential election and the UK referendum on EU membership in 2016, a set of machine learning algorithms were set to work in clustering the features of voter publics. Deep neural networks extracted the features from voters’ data, generating the clusters and patterns that came to constitute and to represent social groups in such a way that they could be microtargeted for political advertising and social media.¹ ‘By profiling every citizen in a country, imputing their personalities and unique behaviours’, writes the former Cambridge Analytica head of research Christopher Wylie, and ‘placing those profiles in a simulation of that society’, the algorithmic model would ‘build the first prototype of the artificial society’.² In fact,

¹A ‘feature’ in computer science is more than simply a characteristic or a property, and though it is often used interchangeably with ‘variable’ it is distinct because it could not be understood in the linear terms of independent or dependent variables. The features of machine learning are more malleable and adjustable than variables and, significantly, they are not defined in advance of the operation. ‘Deep learning takes a different approach to feature design’, explains Kelleher, ‘by attempting to automatically learn the features that are most useful for the task from the raw data’. John D. Kelleher, *Deep Learning* (Cambridge, MA: MIT Press, 2019), p. 32. The different approach to features matters for social science because positivist models have imagined categories of race, class, gender, for example, as though they were measurable variables within a dataset. With machine learning, the features are derived from the representations learned from the dataset.

²Christopher Wylie, *Mindf*ck: Inside Cambridge Analytica’s Plot to Break the World* (New York, NY: Profile Books, 2019), pp. 68–9.

the 2016 machine learning models of UK and US societies were not the first prototype at all. In a twenty-first century variant of the ‘boomerang effect of colonial practice’, Cambridge Analytica’s parent company, SCL Group, had previously modelled societal attributes – creating fractures in communities and harnessing the data by ‘creating havoc’ and ‘riling up crowds’ – in political campaign projects in the Middle East, North Africa, and the Caribbean.³

What could it mean for machine learning technologies to generate a prototypical model of society? Among the many political interventions rendered possible by Cambridge Analytica’s deep neural networks, the African-American communities of many US cities were subject to a campaign of mass voter suppression. Based on models built from multiple data sources, and the clustering of their inferred behaviours and propensities, black voters were microtargeted for anti-Clinton ‘deterrence’ messaging via social media in order to persuade them not to exercise their democratic rights and go out and vote.⁴ The input lines to the machine learning model included voter databases, social media data, and a range of public and commercial datasets from which features and profiles could be extracted. Though the 2016 machine learning models became high profile public cases of disruptive technologies undermining democratic politics, the processes of feature extraction, clustering, and the inference of attributes have become a mainstay technique for the governing of contemporary societies.⁵ From the modelling of social interactions in the COVID-19 pandemic, to the automated streaming of immigration decisions by algorithm, the design of deep learning models has become a form of epistemic politics: a mode of assembling and ordering knowledge of society that fundamentally transforms how state and society comes to understand itself.

Understood as political episteme or order, it is not the case that machine learning technologies merely disrupt the otherwise settled societal norms of good governance, liberal government, or institutional international orders. Indeed, the notion that machine learning algorithms could be subject to good governance via regulation, or ‘AI ethics’, appeals to a different epistemic order than that which is itself generated by deep learning algorithms. Nor should our consideration of the politics of machine learning be confined to the instrumental *application* of AI technologies to specific political domains, and the *implications* for society. Of greater significance for the purposes of this article, the advent of deep learning is generative of new norms and thresholds of what ‘good’, ‘normal’, and ‘stable’ orders look like in the world. That is to say, it is not merely the case that machine learning technologies are supplying new instruments and apparatuses of classification or taxonomy for the governing of society. A significant set of epistemic and political transformations happen when states and societies begin to understand themselves and their problems through the lens of deep neural network algorithms. A machine learning political order

³As Michel Foucault writes, ‘while colonization, with its techniques and its political and juridical weapons, obviously transported European models to other continents, it also had a considerable boomerang effect on the mechanisms of power in the West, and on the apparatuses, institutions and techniques of power. A whole series of colonial models was brought back to the West, and the result was that the West could practice something resembling colonization, or an internal colonialism, on itself.’ Michel Foucault, *Society Must be Defended: Lectures at the Collège de France* (London, UK: Penguin, 2003), p. 103. Recalling SCL’s contract with the Egyptian government in 2013, Christopher Wylie describes the extraction of social media and messaging apps data, and the targeting of social movements ‘to create havoc within the movement’ and ‘riling up crowds’. Wylie, *Mindf*ck*, p. 55. This technique of creating turbulence and fractures in movements, and then harnessing the havoc for political ends – tested and refined in Egypt and Trinidad – was incorporated into the internal colonial projects of the Trump campaign and Vote Leave.

⁴Channel 4 News conducted special investigations into the actions of Cambridge Analytica and SCL, accessing the databases that had been built in the 2016 campaigns, available at: {<https://www.channel4.com/news/revealed-trump-campaign-strategy-to-deter-millions-of-black-americans-from-voting-in-2016>}.

⁵The defining character of machine learning as a computational method is that it has the capacity to learn things that exceed explicitly programmed rules. What this means is that machine learning is a generative process that creates knowledge from the patterns and functions available in data. Machine learning algorithms are thus defined in large part by their iterative relationships to the ‘examples’ they are exposed to in a world of data. From these examples, machine learning extracts the ‘features’ or attributes that are associated with the data example.

does not merely change the political technologies for governing state and society, but is *itself* a reordering of that politics, of what the political can be. What happens, for example, to ideas of political community as a grouping or gathering of people when they become reconstituted as a grouping or clustering of machine-generated features?⁶ What does it mean to infer the attributes of a cluster, and how are categories of race, gender, class, or sexuality understood differently as attributes?⁷ What is at stake for politics itself to be rendered a problem of the design of a model, when every political problem is arranged as a machine learning problem?

In this article I map the contours of machine learning logics as *political orders*. This is not equivalent to the claims that the technologies of deep learning algorithms are becoming political decision-makers, or that there is an automation of previously human bureaucratic processes. Instead, I am concerned with machine learning as devising new limits and thresholds of possibility of how a political project can come into being, what states can do, what a society can be in the future. I have elsewhere detailed a ‘double political foreclosure’ enacted by algorithms, in which ‘the condensing of multiple potentials to a single output that appears to resolve political duress; and the actual preemptive closure of political claims based on data attributes that seek recognizability in advance’ takes place.⁸ In this article, I probe the relationship between the algorithmic foreclosure of political futures and the active programme of the fracturing of communities, the destabilising of democratic rights, and the refiguring of international and social orders.

In many ways the political ordering logics of machine learning run entirely against the grain of notions of shared forms of partial knowledge or, as Elizabeth Povinelli captures ‘alternative projects of embodied sociality’ that have potential for something new to emerge.⁹ When algorithmic systems reduce the intractability and pluridimensionality of politics to a machine learning model, they foreclose the potentiality of other claims and alternative projects that could be built. More than this, they actively incorporate the turbulence, uncertainties, irregularities, and anomalies of fraying and fracturing social relations precisely in order to learn and to modify the model. In short, a machine learning political order is one that profits from the volatilities of fractured disorder. As one of the leading computer scientists of machine learning, Geoffrey Hinton, notes when describing the process of building a model, ‘to capture the variations *we need to learn the features* that it is composed of, and look at the arrangements of those features.’¹⁰ Machine learning algorithms learn to recognise features in the environment through their exposure to variability and contingency, and this process of learning via unknown volatilities is actively enhanced by the breaking and fracturing of social relationships. As witnessed in the wake of the xenophobic targeted media of the UK Vote Leave campaign, for example, the ensuing online

⁶Notwithstanding the performative and constitutive work involved in all forms of the making of social and political groups, my point here is that machine learning is generative of novel forms of grouping, understood as a cluster of attributes. Where Judith Butler argues that a political gathering ‘signifies in excess of what is said’ so that ‘when bodies assemble on the street, in the square ... they are exercising a plural and performative right to appear’, the algorithmic constitution of a gathering reduces this excess and plurality to the single actionable output. Judith Butler, *Notes Toward a Performative Theory of Assembly* (Cambridge, MA: Harvard University Press, 2015), p. 11. In this sense, machine learning political orders both belong to and depart from the inescapable constitutive character of all groups, all gatherings. I am grateful to Martin Coward for challenging my thinking on this question of the constitutive power of grouping and collectivity.

⁷With the use of ‘attribute’ I am foregrounding the slippage between computational ideas of the properties of a cluster of features in data, and political notions of what can be attributed to a person or to a social group. For example, machine learning may extract and cluster the features of a specific group of voters, but it is a political attributive logic that defines the behaviours or propensities of the cluster.

⁸Louise Amoore, *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others* (Durham, NC: Duke University Press, 2020), p. 20.

⁹Elizabeth Povinelli, *Economics of Abandonment: Social Belonging and Endurance in Late Liberalism* (Durham, NC: Duke University Press, 2011), p. 10.

¹⁰The computer scientist Geoff Hinton has been involved in many of the major breakthroughs in machine learning since the 1980s. His most recent research is focused on how to build models that learn functions and extract features without the interventions of a programmer. Geoffrey Hinton, ‘Where do features come from?’, *Cognitive Science*, 38:6 (2014), pp. 1078–101.

racist hate and abuse is itself a useful violent data stream for the refining of attributes of a cluster that is classified as ‘susceptible’ to racist images and messaging. In machine learning political orders, at the very same moment that anti-fascist protest movements are detected by the police’s deep neural networks, another set of algorithms harness the volatility of xenophobia and deepen the risks borne by those bodies who would gather in public space.

In the sections that follow, I am concerned with specifying how a machine learning political order marks some significant discontinuities with traditional statistical and probabilistic imaginaries of state and society. This is only possible through a recognition of how these transformations underwrite the deeply *undisrupted continuities* of racialised violence, police, and state abuses of power, inequality, injustice, and discrimination.¹¹ “Black reason” turns our attention to the technologies (laws, regulations, rituals) that are deployed’, writes Achille Mbembe, ‘as well as the devices that are put in place’ to subject the incalculable to measurement.¹² Machine learning algorithms do indeed subject the incalculable to forms of ordering that deepen the racialised and gendered inscription of violence. What can be arranged of the world and what can be said about it, and by whom? I begin by situating machine learning technologies in a long-standing context of calculative technologies placed at the disposal of the state. If our contemporary world is being shaped by an emergent machine learning political order, then what are its distinct contours? I then discuss three of the principle modes through which the algorithmic state and society begins to understand itself and its problems differently: retroactive design; the indefinite trial; and the foreclosure of spaces of resistance.

Of rules and functions

Among the most significant aspects of the rationale of machine learning, there has been a transformation in what ‘rules’ are in relation to models of the world. Put simply, contemporary machine learning exceeds the programmed ‘if, then, else’ rules of algorithmic decision procedures, seeking instead to generate potential rules and connections from the patterns in the data examples. It is worth spending a little time reflecting on this distinction between rules-based algorithms and deep neural network algorithms. A commonly asked question among observers of emerging non-rules-based deep learning technologies and ‘artificial intelligence’ is what defines deep learning, and how is it novel in the context of extended genealogies of computational rules and cybernetics that have long been intertwined with state power.¹³ The idea of depth in contemporary machine learning refers to the ‘deep’ multilayered neurons that afford the algorithm a capacity to learn representations of input and output data across multiple layers of abstraction. As a group of the world’s most influential computer scientists explain the shift from more traditional approaches of human-engineered rules to deep learning approaches:

Conventional machine-learning techniques were limited in their ability to process natural data in their raw form. For decades, constructing a pattern-recognition or machine-learning system required careful engineering and considerable domain expertise ... Deep learning methods are representation-learning methods with multiple levels of representation, obtained by composing non-linear models that each transform the representation at one level into a representation at a higher, slightly more abstract level. With the composition of enough such transformations, very complex functions can be learned. An image, for example, comes in the form of an array of pixel values, and the learned features in the first layer of

¹¹Simone Browne, *Dark Matters: On the Surveillance of Blackness* (Durham, NC: Duke University Press, 2015); Ruha Benjamin, *Race After Technology: Abolitionist Tools for the New Jim Code* (Cambridge, UK: Polity, 2019).

¹²Achille Mbembe, *Critique of Black Reason* (Durham, NC: Duke University Press, 2017), p. 31.

¹³Orit Halpern, *Beautiful Data: A History of Vision and Reason since 1945* (Durham, NC: Duke University Press, 2014); N. Katherine Hayles, *My Mother Was a Computer: Digital Subjects and Literary Texts* (Chicago, IL: Chicago University Press, 2005).

the representation typically represent the presence or absence of edges at particular locations in the image. The second layer typically detects motifs by spotting particular arrangements of edges, regardless of small variations in the edge positions. The third layer may assemble motifs into larger combinations that correspond to parts of familiar objects, and subsequent layers would detect objects as combinations of these parts. The key aspect of deep learning is that *these layers of features are not designed by human engineers*: they are learned from data using a general-purpose learning procedure.¹⁴

To put the distinction simply, where past approaches to machine learning began with the engineering of domain-specific rules that captured the relationship of variables within a dataset, deep learning is understood to decompose a problem into multiple layers of representation of patterns in the data (often called ‘hidden layers’). In this specific sense, deep learning algorithms are more experimental and open-ended in their computation than past rules-based forms. Indeed, the computer scientists make a direct association between: a world of more complex multidimensional issues to address; a greater abundance of available ‘big’ data; and the perceived limitations of human-engineered rules-based algorithms. For example, as Geoff Hinton describes the difference between rules and learned features, ‘there are some problems where it is very hard to write the program, there may not be any rules’ so that ‘instead of writing a program by hand for each specific task, a machine learning approach collects lots of examples, finds clusters in the input, and provides a representation of the input in terms of learned features’.¹⁵ Here the engineering problem itself is differently understood, with the difficulty of the task (a twinned difficulty of a complex world and a computational task with high dimensionality of data) allied to the detection of features *not designed by human engineers*. Here one begins to see the notion that there are problems so complex, so multiple in their dimensions, that it is no longer possible for a human engineer to determine the variables, define the rules, and write the program.

There is a significant point of resonance, then, between computational and political ideas about rules under conditions of complexity and interdependence.¹⁶ To be clear, this is not a causal relationship where ideas from computer science bleed into the state and sovereign logics, but a correlative one in which the state acts as a kind of ‘resonance chamber’ that entangles and amplifies the power of deep learning models.¹⁷ Specifically, there is a resonance between the idea that complex political problems may exceed conventional bodies of knowledge or traditional statistical and probabilistic epistemes, and the computational idea that the features of a machine learning model are not knowable in advance. The unravelling and fraying of attachments to postwar political models of the welfare state and liberal international norms, I suggest, is entangled with the undercutting of rules-based algorithmic systems.¹⁸

What does it mean for rules-based algorithms to be undercut by deep learning? If one considers algorithmic rules to be ‘if ... and ... then ... else’ formulations, then a rules-based model for credit risk, for instance, arranges propositions such as: IF income < X AND savings < Y THEN high-risk ELSE low-risk.¹⁹ It is precisely such rules-based logics that characterised

¹⁴Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, ‘Deep learning’, *Nature*, 521:28 (2015), p. 436, emphasis added.

¹⁵Geoffrey Hinton, ‘Neural Networks for Machine Learning’, Coursera lecture (2019), available at: {<https://www.cs.toronto.edu/~hinton/coursera/lecture1/lec1.pdf>} accessed September 2021.

¹⁶Mark C. Taylor, *The Moment of Complexity: Emerging Network Culture* (Chicago, IL: Chicago University Press, 2001).

¹⁷The state is not a point taking all the others upon itself, write Deleuze and Guattari, ‘but a resonance chamber for them all’. The ‘entanglement of the lines’ of power is thus not a linear causal relation but a form of amplification that extends the reach of the state. Gilles Deleuze and Felix Guattari, *A Thousand Plateaus* (London, UK: Continuum, 2004), p. 247.

¹⁸I use attachment here to signal more than a commitment to social democratic institutions or welfare state models, and to conceive of attachments as Lauren Berlant understands them, as ‘structures of relationality’ in which one invests hope for specific futures. Lauren Berlant, *Cruel Optimism* (Durham, NC: Duke University Press, 2011), p. 13.

¹⁹Ethem Alpaydin, *Machine Learning* (Cambridge, MA: MIT Press, 2016), p. 49. See also Taina Bucher, *If ... Then: Algorithmic Power and Politics* (Oxford, UK: Oxford University Press, 2018).

the enrolment of data-mining technologies into the war on terror, so that computer scientist Rakesh Agrawal declared of his homeland security models, ‘this is a site in Pakistan and what is happening here is these different characteristics we are interested in, they are written into the rules’.²⁰ The engineering of algorithmic systems for the state was thus broadly locatable in a human engineer, and usually also a knowledge domain specialist, who would define the characteristics of interest and encode them into programmable rules. The racialised logics of the characteristics of interest in Pakistani online images and text – as seen in Agrawal’s engineered rules – closely allied settler colonial state enmities with mathematical rule building. The association rules of Agrawal’s models exhibit what Achille Mbembe has described as ‘the process by which certain spaces are transformed into uncrossable places for certain classes of populations, who thereby undergo a process of racialization’.²¹ In short, the engineering of rules is the engineering of race into a probabilistic system that classifies people according to degrees of riskiness.

For much of the second half of the twentieth century, and into the first decade of the twenty-first, it remained sequential rules-based algorithms – and their attendant racialised statistical probabilistic forms of knowledge – that shaped the practices of the political and policy sciences, public administration, and foreign policy. The dominant rational models of politics and policy sought to build stability and order amid the turbulence and uncertainties of the cold war. Computational logics were an important component of postwar rationality and rules-based orders precisely because they extended the sequential rules of the algorithm into politics, administration, and decision.²² In the context of uncertainty and nuclear threat, the engineering of rules rendered decision pathways as sequential steps that would optimise outcomes and guard against capricious human actions:

Although they often described their epoch in terms of complexity, uncertainty, and risk, and conjured the spectre of a nuclear war triggered by accident, misunderstanding, or lunacy, the participants in the debate over Cold War rationality believed that the crystalline definiteness, generality, and conclusiveness of the algorithm could cope with a world on the brink ... *Algorithmic rules* impervious to context and immune to discretion, rules that could be executed by any computer, human or otherwise ... came together as a new form of rationality with glittering cachet in the human sciences and beyond.²³

Understood as definitive and conclusive algorithmic procedures, the idea of rules that could be executed by any computer – human or machine – substantially underwrote the architectures of postwar social and international orders.²⁴ From the New Deal rules and regulatory arrangements of the US, to the building of welfare state rules and entitlements in Europe, the growth of new functions of government required computational models that were rules-based, invariant, and

²⁰Rakesh Agrawal, *The Mathematical Sciences’ Role in Homeland Security* (Washington, DC: BMSA and National Research Council, 2004). For discussion of how data mining wrote characteristics into rules, see Louise Amoore, *The Politics of Possibility: Risk and Security Beyond Probability* (Durham, NC: Duke University Press, 2013), and Marieke de Goede, *Speculative Security: The Politics of Pursuing Terrorist Monies* (Minneapolis, MN: University of Minnesota Press, 2012).

²¹Achille Mbembe, ‘Bodies as borders’, *From the European South* (2019), pp. 5–18 (pp. 4, 9).

²²Herbert Simon’s *Models of Man* extended mathematical formulae into human behaviour, proposing the ‘bounded rationality’ within which individuals build simplified models of complex situations. Herbert Simon, *Models of Man: Social and Rational: Mathematical Essays on Rational Human Behaviour in a Social Setting* (New York, NY: Wiley, 1957), p. 203. John von Neumann and Oskar Morgenstern established the major mathematical theory of social and economic organisation. Von Neumann’s later work would model human behaviour as a linear programming problem. John Von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behaviour* (Princeton, NJ: Princeton University Press, 1944).

²³Paul J. Erickson, Judy Klein, Lorraine Daston, Rebecca Lemov, Thomas Sturm, and Michael Gordin, *How Reason Almost Lost Its Mind: The Strange Career of Cold War Reasoning* (Chicago, IL: Chicago University Press, 2013), p. 31, emphasis added.

²⁴Robert MacBride, *The Automated State: Computer Systems as a New Force in Society* (New York, NY: Chilton Book Co., 1967).

replicable in all circumstances. A pivotal part of this twinned alignment of political and computational rule building is the concept of *function*. The growth of neo-functionalism and regime theory in international studies, for example, established a particular meaning of ‘function’ that owes distinct debts to cybernetic theory and cybernetic models of computation. Consider Ernst Haas’s founding mid-twentieth-century text of neo-functionalism, in which he details the institutional building of the European Coal and Steel Community.²⁵ As Ben Rosamond proposes in his reading of Haas, neo-functionalism sought ‘a set of general propositions’ and ‘a set of candidate independent variables’ from which the ‘likelihood’ of political outcomes could be derived.²⁶ At the heart of neo-functionalist accounts of political community is the idea that actions in the name of one function or policy domain will ‘spillover’ into new functions, bringing new momentum for political integration. Such a logic of the overflow from one domain to the next, Rosamond argues, is ‘suggestive of automaticity’ in the sense that the functions of parts of a system will, in aggregate, be somehow exceeded and overtaken by the functioning of a whole.²⁷ Indeed, in Haas’s later writings he himself acknowledges the place of the cybernetic thought of Karl Deutsch in his formulation of political knowledges, proposing that ‘in cybernetic terms, the strategy calls for more effective information transfer among actors and organizers to enhance survival potential.’²⁸ From the neo-functionalist perspective, the probability of particular desired political outcomes (for example, interdependent relations among nations, integrated functions in international organisations) could be sought through independent variables and general propositions.

The specific postwar orders of political and computational rules, functions, and variables are of significance to how one understands the emergence of contemporary machine learning as a political order. It is my argument that the transformation from rules-based algorithms to deep learning models has also been a condition of possibility for the undoing of rules-based social and international orders, from the Brexit challenges to EU integration, to the austerity politics and digitalisation of welfare states. Where rules-based computation and decision was critical to the formation of postwar international liberal order, and to the formation of welfare states in twentieth century, what happens when the machine learning function displaces it? Though the cybernetic mathematical and military sciences are important origins of contemporary deep learning technologies, there is also a significant transformation to the idea that algorithmic rules extend into human decision-making. In a sense, machine learning’s *raison d’être* is to generate outputs that are in excess of the formulation of rules, something not determinable in advance.²⁹ What we are witnessing, in short, is a transformation from algorithmic rules conceived to tame a turbulent, divided, and capricious world, to the productive generation of turbulence and division from which algorithmic functions are derived.

The process of machine learning finds an optimal ‘function’ by mapping the representations of input data in order to achieve a target output. So, because computer science proposes that all existing functions can be approximated by a neural network, the algorithmic political arrangement becomes one in which all political problems can be figured as machine learning problems. Consider, for example, how a political question becomes refigured in and through the propositions of machine learning: ‘what is the optimal representation of all the input immigration data to achieve this target of limited immigration?’; ‘what is the best representation of all human mobility data to achieve the target of limiting COVID-19 transmission?’; ‘what is the representation of crime data that optimizes the output of urban policing in this district of

²⁵Ernst Haas, *The Uniting of Europe: Political, Social, and Economic Forces, 1950–1957* (Notre Dame, IL: University of Notre Dame Press, 1958).

²⁶Ben Rosamond, ‘The uniting of Europe and the foundation of EU studies: Revisiting the neofunctionalism of Ernst B. Haas’, *Journal of European Public Policy*, 12:2 (2005), pp. 239, 246.

²⁷*Ibid.*, p. 244.

²⁸Ernst Haas, ‘Is there a hole in the whole? Knowledge, technology, interdependence, and the construction of international regimes’, *International Organization*, 29:3 (1975), p. 845.

²⁹Luciana Parisi, *Contagious Architecture: Computation, Aesthetics, and Space* (Cambridge, MA: MIT Press, 2013).

London?'. It is for this reason that it is insufficient to merely say that automated technologies or machine learning systems disrupt our social order, or undercut our existing bodies of rights. It is more significant, even, than this disruptive force. For it is itself a mode of politics that arranges the orderings of public space, adjudicates what a claimable right could be, discriminates the bodies of those on whom it is enacted.

Why does it matter that algorithms are now predominantly generating their own rules from the contingencies of data, or that they derive functions from representations of data? What is at stake? In common with derivative financial and digital formations that have the capacity to unbundle and trade attributes indifferent to the underlying asset values as Randy Martin's work details³⁰ – the neural network computes functions of input and output data indifferent to the underlying context of that data; rendering the derivative function tradeable and exchangeable across domains.³¹ The idea of a function thus becomes unmoored and disconnected from underlying values so that what is prized is its very flexibility and adaptability to shocks and disruptions across domains. As Orit Halpern writes on the logic of resilience, it 'has a peculiar logic. It is not about a future that is better, but rather about an ecology that can absorb constant shocks while maintaining its functionality and organization.'³² Machine learning logics embrace this resilience thinking in the way that they continually modify via their exposure to the surprises of new data. In this way, algorithms are no longer supplying the stabilising sequences of rules to an otherwise fraught state bureaucracy, but are precisely decoupling data from any sense of an underlying function.

Where one might envisage adjudicating the success of a policy decision on the basis of whether it has achieved a stated function, the machine learning model can always approximate a function and is, therefore, indifferent to success or failure as such. As Martin writes, 'compared to earlier imperial forms, the empire of indifference stands as a massive flight from commitment, urging an embrace of risk and self-management' then 'ignoring, incarcerating, or dispossessing those who cannot make the grade'.³³ The indifference to underlying values that Martin so powerfully depicts in relation to war and finance is present also in the vast expansion of deep learning models that ignore, incarcerate, or dispossess those who are not a good fit to the model. The model's ability to trade volatilities, indifferent to the consequences – as Orit Halpern and Randy Martin differently propose – is not about committing to a better future, but rather it valued precisely for its capacity to absorb external shocks and surprises, come what may.

In sum, whereas sequential and rules-based algorithms resonated with rules-based political orders, with equilibrium and stability considered achievable via the functional and mechanical arrangement of inputs and outputs, the rise of non-rules-based generative algorithms deploys clusters and patterns of attributes to gear up for disequilibrium and volatility. Deep learning political orders detach their derivative attributes from underlying social values in the broadest sense. A domain agnostic and free-floating function is no longer tethered to conventions of good governance or a good society, but is instead actively generating new notions of what good, bad, normal, or abnormal could be in a society.

³⁰Randy Martin, *An Empire of Indifference: American War and the Financial Logic of Risk Management* (Durham, NC: Duke University Press, 2007).

³¹'The central characteristic of derivatives', write Dick Bryan and Michael Rafferty, 'is their capacity to dismantle or unbundle any asset into constituent attributes and trade those attributes without trading the asset itself.' Richard Bryan and Michael Rafferty, *Capitalism with Derivatives* (Basingstoke, UK: Macmillan, 2006), p. 44. It is the trading of attributes without the underlying asset itself that is present also in what I have elsewhere called 'data derivatives' that unbundle and trade attributes indifferent to underlying values. Louise Amoore, 'Data derivatives: On the emergence of a security risk calculus for our times', *Theory, Culture and Society*, 28:6 (2011), pp. 24–43.

³²Orit Halpern, 'Hopeful resilience', *e-flux Architecture* (19 April 2017), p. 4. See also Kevin Grove, *Resilience* (New York, NY: Routledge, 2018).

³³Martin, *An Empire of Indifference*, p. 14.

Retroactive design: How the solution gets the problem it deserves

What does it mean to say that a political problem becomes configured as a machine learning problem? The advent of deep learning algorithms – and their penetration into many aspects of life – has witnessed the process of modelling a society becoming a political end in itself. A significant defining aspect of the process of becoming a political end is the retroactive logic of beginning with an end target and abductively working back to adjust the parameters of the model in order to converge on the target.³⁴ The abductive logics of machine learning algorithms, as Luciana Parisi proposes, ‘construct a new kind of model, which derives its rules from contingencies and open-ended solutions’.³⁵ The end target of machine learning could be anything. It could be an immigration target from which all the parameters of immigration are to be abductively modelled; a distribution of grades in a pandemic exams algorithm; or it could be a threshold target of national health service capacity from which pandemic rules are to be abductively generated. In such policy cases, the setting of the target output – such as an immigration ‘cap’, or a parameter of NHS intensive care capacity – becomes the constrained limit of what is politically possible. The point is that the end target as starting point significantly reconfigures the relationship between the formulation of a political problem and the proposition of a solution.

This notion of beginning from a solution and working back to the problem requires some further elaboration. ‘The problem always has the solution it deserves’, writes Gilles Deleuze, so that a solution is constituted by ‘the way in which it is stated’ and ‘the conditions under which it is determined as a problem’.³⁶ Though there may be multiple potential solutions proposed to a political problem – what should a border do? How should the health of the population be secured, and how should it be funded? How is unemployment to be managed? – the plural potential solutions are nonetheless unified across their differences by the conditions under which the problem is determined. It is precisely this problem-solution relation that also characterises Foucault’s use of ‘problematization’ to capture ‘a domain of acts, practices, and thoughts’ that generate the ‘conditions in which possible responses can be given’.³⁷ As Foucault explains in an interview during the last year of his life, across his historical thought on madness, penal codes, and sexuality, what came to matter was to ‘rediscover at the root of these diverse solutions the general form of problematization that has made them possible’.³⁸ Such methods of tracing the genealogy of a solution to political problems have animated analysis of how sociotechnical solutions have become political technologies.³⁹ Here, the solution is a function of how the problem was framed, so that the problematization comes prior to the solution and explains its ontology. In the terms of algorithmic rules, the problematization is akin to the conditional ‘if’ in an ‘if ... then’ formulation. For much of the second half of the twentieth century, this relationship between the problematization of stability, state security, and reconstruction, and the diverse solutions of welfare states, public health and housing, and liberal international organisation, endured and persisted across its differences.

Yet, what I propose here is that machine learning political orders reverse Deleuze’s dictum so that the political problem is constituted by the posited solution; or, *the solution gets the problem it*

³⁴J. Josephson and S. Josephson, *Abductive Inference: Computation, Philosophy, Technology* (Cambridge, UK: Cambridge University Press, 1996).

³⁵Parisi, *Contagious Architecture*, p. 2. On abductive logics of algorithms, see also Claudia Aradau and Tobias Blanke, ‘Governing others: Anomaly and the algorithmic subject of security’, *European Journal of International Security*, 3:1 (2017), pp. 1–21; and Louise Amoore and Rita Raley, ‘Securing with algorithms: Knowledge, decision, sovereignty’, *Security Dialogue*, 48:1 (2017), p. 6.

³⁶Gilles Deleuze, *Bergsonism* (New York, NY: Zone Books, 1991), p. 16.

³⁷Michel Foucault (with Paul Rabinow), ‘Polemics, politics, and problematizations: An interview with Michel Foucault’, in Paul Rabinow and Nikolas Rose (eds), *The Essential Foucault* (New York, NY: The New Press, 2003), pp. 20, 24.

³⁸*Ibid.*, p. 24.

³⁹Francois Ewald, *L’Etat Providence* (Paris: Bernard Grasset, 1986); Michael Dillon, *Politics of Security: Towards a Political Philosophy of Continental Thought* (London, UK: Routledge, 1996).

deserves. The retroactive logic of deep learning commonly begins with the identification of the target output from the model, actively using the output signals that diverge or converge on the target as an experimental space of modulation and adjustment. Though there are multiple possible functions or ‘solutions that will match the data’, a machine learning algorithm will use ‘two sources of information to select the best function: one is the dataset, and the other (the inductive bias) is the assumptions that bias the algorithm to prefer some functions over others’.⁴⁰ In order to change the output signal, the weights within the hidden layers of the neural network adjust and modify the signal that feeds forward to the next layer. Why does this computational process matter for how the political problem and the solution interact? The retroactive move from target solution to the weights in the model means that the parameters and dimensions of an intractably difficult political question – democracy, pandemic response, border security, stability in the economy – become configured as infinitely adjustable in relation to the solution. Where the concept of problematisation suggests a multiplicity of actions that could take place under the broad unifying conditions of the formulation of the problem – and, indeed a space for normative deliberation of possible actions – this retroactive paradigm forecloses the multiplicity of plural solutions to a single target, and reduces the framing of the political problem to the weighting of inputs. Every adjustment or modification of the parameters in the deep learning model is simultaneously an arrangement of the political problem.

To begin at the ‘end’, with a target output of the model, is thus to transform radically what a political claim can be in the world. For example, machine learning algorithms are increasingly being deployed in immigration and borders decisions. The introduction of machine learning in the sorting and classification of visa applications does not merely automate some aspects of a previously human centred bureaucratic process. More than this, the building of a model of the flows of immigration claims actively constitutes what a border can be and how it is understood as a political question. Moreover, the space to challenge the political formulation, such as for example to question the racialised criteria that are applied to the judgement on a person, is also closed out by the machine learning process. In contrast to the rules-based models that I described as engineering racialised assumptions into association rules (for example, in Rakesh Agrawal’s data mining models), machine learning learns and generates new racial formations from the data examples. With machine learning forms of classification, there are no criteria as such; there are only inputs, outputs, features, and functions. The model will also adjust with the volatilities and geopolitical disruptions of migration, shifting its thresholds so that the derived outputs are decoupled from the input data of a visa application.

There are profound political consequences of this more generative and disruptive approach to data inputs. In 2020, for example, the Joint Council for the Welfare of Immigrants (JCWI) challenged the UK Home Office’s use of a ‘streaming algorithm’ to allocate visa applications to high, medium, and low levels of risk. The JCWI identified the nationality data that were among the inputs to the algorithm that scored some applicants as ‘high risk’, effectively automating the decision to refuse the visa. JCWI and Foxglove Legal successfully argued that the nationality data are proxies for race and, therefore, in breach of the provisions of the Equality Act 2010.⁴¹ If the algorithm had been a solely rules-based ‘if ... and ... then ... else’ sequence, then the juridical removal of racist input data would arguably substantially change the outputs of the system. However, with machine learning it would be a mistake to conclude that the removal of a racist input will excise the racialised propositions of the model. The streaming of visa applications into risk-rated clusters, as exhibited in the JCWI case brought by Foxglove, is an example of a solution defining and configuring the political problem, so that immigration targets are the

⁴⁰Kelleher, *Deep Learning*, p. 18.

⁴¹*JCVI v. Secretary of State for the Home Department* (2020), full papers available at: {<https://www.foxglove.org.uk/news/home-office-says-it-will-abandon-its-racist-visa-algorithm-nbsp-after-we-sued-them>} accessed September 2021.

starting point. When the input data are not *variables* (in the functionalist sense) but *features*, the UK Home Office can agree to remove a racist input (as they have done) while continuing to weight other features in ways that revitalise racist inferences that were not strictly present in the input data. For example, the weighting and parameters applied to travel, to familial relationships, or to periods of time spent outside the UK can serve to constitute a suspicious population and to generate a red-flagged output. When the model is learning about salient features and clusters from the dataset, its racialised assumptions will exceed the categories of the input data and extend to the groupings and communities created by the machine learning process. It is not only the use of data from which race can be inferred, but more significantly that the immigration algorithm forecloses the potential of a person's future on the basis of a racism that pervades the model all the way down.⁴² In short, the question of what kinds of political actions, which political claims, or which policy agendas can be designed and made, is condensed down to the foreclosure of a target output solution.

When the design of a machine learning model becomes a valued political object in itself, the derivative outputs of the model are exchangeable and tradeable beyond any specific defined political problem or 'domain'.⁴³ In common with the models built by Cambridge Analytica, everything becomes a function of deep learning to the point that there are no bad outputs – even where this may be illness, racism, death, destitution, social hardship, child poverty – there are only target outputs and the adjustable parameters of a problem. Amid the loss of more than 170,000 lives to COVID-19, the UK launched its national data strategy in 2020, describing the 'high watermark of data use set during the pandemic' where businesses, government, and organisations had been 'freed up' to 'innovate and experiment'.⁴⁴

A period marked by lack of effective emergency planning and horrifying loss of life is thus articulated as a 'high watermark of data use'. In his testimony before the UK House of Commons Science and Technology Select Committee, former advisor to the prime minister, Dominic Cummings, explained that conventional civil contingencies 'did not have the data architecture', and that 'companies [DeepMind] had stuff we could use off the shelf, hack it together for the NHS'.⁴⁵ Behind the rhetoric of his testimony, Cummings's account does express the deep faith placed in deep learning models to address the data gaps in bureaucratic structures. On the day that he appeared before MPs, Cummings released a photograph of a whiteboard used to map early pandemic planning. Among the scrawled notes, a question is posed: 'who do we not save?'. Viewed from the perspective of the building of machine learning models for the pandemic – or the 'hacking together' by private tech companies – the question of who will not be saved is but a mere parameter in a model, to be adjusted in relation to NHS capacity. In the event, this parameter was also a deeply racialised metric – with people of black and South Asian ethnic background four times more likely to die from COVID-19 in the UK.⁴⁶ In common with the racialised logics of the immigration algorithm, the pandemic models did not need to begin with race as a category or input to nonetheless generate a deeply racialised model of algorithmic violence. The machine learning model itself has extraordinary resilience in the face of complete moral and political failures because a weight can always be adjusted, a

⁴²Ruha Benjamin, 'Introduction: Discriminatory design, liberating imagination', in Ruha Benjamin (ed.), *Captivating Technology: Race, Carceral Technoscience, and Liberatory Imagination in Everyday Life* (Durham, NC: Duke University Press, 2019).

⁴³David Ribes, Andrew Hoffman, Steven Slota, and Geoffrey Bowker, 'The logic of domains', *Social Studies of Science*, 49:3 (2019), pp. 281–309.

⁴⁴UK National Data Strategy (2020), available at: {<https://www.gov.uk/government/publications/uk-national-data-strategy/national-data-strategy>} accessed September 2021.

⁴⁵Dominic Cummings's testimony before the UK Science and Technology Select Committee, transcript available at: {<https://committees.parliament.uk/oralevidence/2249/pdf/>} accessed December 2021.

⁴⁶Robert Booth and Caelainn Barr, 'Black people four times more likely to die from Covid 19, ONS finds', *The Guardian* (7 May 2020).

threshold modified, a parameter tweaked. The question of ‘who do we not save?’ is translated into the parameter of a model whose target outputs are the starting point. In a situation where there is a total and abject failure of policy and good governance, the innovation in data science and AI is nonetheless fostered by the racialised violence and social turbulence that is generated. Whether this is incorrect or unjust decisions made in an algorithmic benefits system, poor judgements on policing deployment, or the catastrophic pursuit of a modelled ‘herd immunity’, machine learning political orders learn something from the data generated by the volatility.

A similar sense of the productivity of fractured governance is present in the UNHCR’s statements that ‘even in highly volatile and chaotic environments’, digital systems will ‘radically expand the responses that can be crafted for challenges in health care, education, migration, and security’.⁴⁷ The organisation envisages machine learning technologies within a process of ‘competitive disruption’ to what it calls ‘obsolete’ institutional structures ‘with legacies dating back to World War II’. The flexibility and agility of a deep learning model – deployed, for example, in UNHCR’s ‘Project Jetson’ predictive models of refugee movements – becomes a condition of possibility for the imagination of adaptive digital humanitarian and pandemic response, so that social and political relations are reconfigured as the parameters of a model.⁴⁸

In each of the examples I describe above, functionally arranged structures of postwar social and international orders are reimagined along the dimensions of a machine learning universal function. That is to say, a function that is immanently mappable from a target output to the weightings of the layers of the problem. To propose that a policy or an institution must deliver on a function thus also shifts its ground – for something to ‘function’ it no longer needs to work as such.⁴⁹ As Debbie Lisle has argued, the cultures of science and engineering mobilise a politics within which ‘failure’ itself is rendered an ‘instructive experience’.⁵⁰ Within a machine learning logic, the instructive experience of failure permits the model to learn those unknown things that are beyond the distribution of data in a training dataset. Though machine learning orders cannot be said to fail as such – or at least the output of the model is never a failure but only a signal – the retroactive generation of political problem from output solution means that the very idea that a neural net can approximate *any function* becomes a powerful political idea. In short, though these ideas of failure are ontologically distinct, they become epistemically aligned; there is slippage between failure as learning, and the idea that there can be no ethical failure, nor catastrophic policy failure.

At stake in retroactive design, then, is not merely that deep learning algorithms are deployed to govern society, but rather that society comes to understand itself and its problems through the lens of the deep learning model. The relations among people, objects, and space become rendered as features from which something useful can always be extracted, from which a function can always be found. However, the plurality and multiplicity of those relations, and the potentiality for new or alternative political projects to emerge, is radically foreclosed around the retroactive mapping from target output to the weighting of parameters in a model. It is to the usefulness of exposure to contingency that I now turn, with a discussion of how technological trials became perennial and indefinite.

⁴⁷UNHCR, ‘Disruption and Digital Revolution for Whom?’, available at: {https://www.unhcr.org/innovation/wp-content/uploads/2020/05/Disruption-and-digital-revolution-for-whom_WEB052020.pdf}.

⁴⁸UNHR, ‘Project Jetson’, available at: {<https://jetson.unhcr.org/>}. On the implications of algorithmic models for humanitarian response, see Mark Duffield, *Post-Humanitarianism: Governing Precarity in the Digital World* (Cambridge, UK: Polity, 2018).

⁴⁹Jacqueline Best, *Governing Failure* (Cambridge, UK: Cambridge University Press, 2014).

⁵⁰Debbie Lisle, ‘Failing worse? Science, security and the birth of a border technology’, *European Journal of International Relations*, 24:4 (2017), pp. 887–910.

Trial by design: How the test became an indefinite trial

The concept of alpha and beta testing has its origins in IBM's 1950s software development, when the 'A' test signalled the in-house testing to improve the engineering and the 'B' test referred to the verification through user engagement and development. In his account of Cold War America and the culture and politics of computers, Paul Edwards describes the emergence of cybernetic culture and its penetration of state and military thinking. In the context of Edwards's 'closed world' integration of a 'seamless web' of human and machine systems, the engineer deploys the testing regime in a form that mirrors the hierarchical logics of computers.⁵¹ As an engineering practice of the twentieth century, the test formed an intrinsic part of the rules-based and cybernetic approaches to government and computer science. As Edwards depicts the Cold War collaborations of IBM, RAND, and the aircraft corporations, the process of political planning itself became an 'if ... then' proposition, 'constructing a list of tests to perform', identifying failings as information problems, and creating feedback loops from test to engineer.⁵² In this sense, the Cold War alliances between computer science, mathematics, and the military state embodied a specific understanding of testing as practice, and of errors as problems of fallible human perception that could be corrected with machine systems.⁵³

This conception of testing and the sequential procedures of the 'list of tests to perform' is aligned with the rules-based computation and rules-based social and international political orders I have described. It is a conception of the test that is deeply probabilistic and conceives as testing as a process concerned with the calculation of probabilities. As Orit Halpern describes the cyberneticians' practices, 'they focused on the ability to calculate the probability that one set of interactions (the missile hitting the plane) will occur, over other sets less likely but possible'. 'This is a worldview composed of functionally similar entities – black boxes', she writes, composed only of 'their algorithmic actions in constant conversation with each other producing a range of probabilistic scenarios'.⁵⁴

Against this historical backdrop of probabilistic alpha and beta testing within functionally similar entities, the rise of deep learning algorithms is most profoundly possibilistic in its orientation to the future.⁵⁵ As a mode of political ordering, machine learning circumvents modern notions of testing in science and engineering by turning to *the trial* and *trailing* as experimental technology. The trial is a more possibilistic approach precisely because it refutes the functionally similar entities and probabilities Halpern denotes in cybernetics, and it embraces instead the generation of multiple possible functions in order to defer a decision on what is politically useful. Understood in this way, the rise of trailing in contemporary machine learning has more in common with the conduct of stress testing to anticipate uncertainty in finance than it does with alpha and beta testing in software engineering.⁵⁶ The machine learning model dwells indefinitely in its trial phase because it is designed and redesigned through its exposure to people, objects, places, and scenes, perennially modifying itself in response to what it has learned through its encounters. In this way, the 'demo' as technological demonstration, has a close relationship with the 'demos' as the people, the population,

⁵¹Paul Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America* (Cambridge, MA: MIT Press, 1996), pp. 1, 232.

⁵²Ibid., p. 232.

⁵³Ibid., pp. 20–1.

⁵⁴Halpern, *Beautiful Data*, p. 46.

⁵⁵I have elsewhere distinguished probabilistic from possibilistic modes of risk and calculative futures. The possibilistic logic 'does not deploy statistical probabilistic calculation in order to avert future risks but rather flourishes in conditions of declared constant emergency because decisions are taken on the basis of future possibilities, however improbable'. Amoores, *The Politics of Possibility*, p. 12.

⁵⁶Paul Langley, 'Anticipating uncertainty, reviving risk? On the stress testing of finance in crisis', *Economy and Society*, 42:1 (2013), pp. 51–73.

and democracy.⁵⁷ ‘Our forms of technological testing and deming’, writes Halpern, ‘envision a world where artificial intelligences and computers can replace the democracy that is now imagined to be obsolescent’.⁵⁸ As deep learning models penetrate public space, for instance in live facial recognition biometric systems in urban spaces, at borders, and in military spaces, every trial of a deep learning model is also an active reconfiguration of that space as the model adapts in response to the contingencies it yields.⁵⁹

For example, in the world’s first legal challenge to police use of automated facial recognition algorithms (AFR), the appellant, Bridges, argued that South Wales Police unlawfully extracted his biometric data during two trials of the technology.⁶⁰ Bridges had been subject to AFR during a protest outside an arms fair in Cardiff in 2018, and during a Christmas shopping trip in 2017, with each trial of the system storing his biometric data, cross-matching with a watchlist, generating match scores, and modifying the sensitivity of the model. The court of appeal found in Bridges favour in 2020, following testimony from an expert computer scientist whose account vividly illustrates how the trial indelibly marks and recalibrates a gendered and racialised system. ‘AFR systems will have a higher error rate for women and people from black and ethnic minority groups’, he testified, and ‘where an end user is adjusting threshold values it may make the AFR system particularly sensitive for some individuals. People from that ethnic group will be wrongly matched more often.’ Thus, the trial of AFR – ongoing for a seemingly indefinite period from 2017 – will continue to generate racialised outputs and clusters of suspicion, even where individual biometric datasets are deleted. As Rocco Bellanova and Marieke de Goede describe architectures of data analysis, ‘the infrastructure aims at defining the “right population” to be algorithmically governed.’⁶¹ The very communities who are already disproportionately targeted by the state will experience an intensification of scrutiny and control. In this way, the capacity of a person to be present, or to gather with others, in public space is iteratively and intimately related to the exposures of a machine learning model that is indefinitely trialled across multiple spaces. Unlike the feedback loops of Edwards’s Cold War military-computer science collaborations, the error rate of the biometrics are contingent on the shifting infrastructural thresholds and parameters of the algorithm. Whereas the cybernetic mode of testing was concerned with the engineering of human and machinic component parts, the machine learning mode of indefinite trials makes the limit and the threshold the object of the trial, so that setting sensitivities, moving borders and boundaries reconfigures both algorithm and action.

In this way, the orientation of the indefinite trial is closer to an experimental and open-ended process of design than it is to engineering. The very etymology of design is from the Latin *designare*, to designate, to mark out, and related to *disegnare*, to contrive or intend. It is precisely this process of designating and marking out that I see at work in the indefinite trials of deep learning technologies in cities, at borders, in public space. Bruno Latour outlines a philosophy of design in which ‘design has been extended from the details of daily objects to cities, landscapes, nations, cultures, bodies, genes.’⁶² For Latour, the practice of design is counterposed to historical notions of building or engineering, so that ‘things are no longer “made” or “fabricated”, but rather

⁵⁷Orit Halpern, ‘Demos’, in Nanna Bonde Thystrup, Daniela Agostinho, Annie Ring, Catherine D’Ignazio, and Kristin Veil (eds), *Uncertain Archives: Critical Keywords for Big Data* (Cambridge, MA: MIT Press, 2021).

⁵⁸Halpern, ‘Demos’, p. 134.

⁵⁹The trialling of machine learning technologies across multiple domains acts to render public spaces as ‘feature spaces’, where the data features of an environment actively shape the algorithms (Amoore, *Cloud Ethics*, p. 58), and with private tech companies such as Palantir, Idemia, Nvidia, and NEC selling their systems in new places precisely on the basis of their trial in another domain.

⁶⁰*Bridges v. South Wales Police*, available at: {<https://www.judiciary.uk/wp-content/uploads/2020/08/R-Bridges-v-CC-South-Wales-ors-Judgment.pdf>} accessed September 2021.

⁶¹Rocco Bellanova and Marieke de Goede, ‘The algorithmic regulation of security: An infrastructural perspective’, *Regulation and Governance*, online (2020), p. 8.

⁶²Bruno Latour, ‘A Cautious Prometheus? A Few Steps Toward a Philosophy of Design’, keynote lecture for the Design History Society, Falmouth, Cornwall (3 September 2008).

designed.’ ‘This was the old way’, he writes, ‘to build, to construct, to destroy, to radically overhaul’ through engineering.⁶³ By contrast, to design something, for Latour, is never to found something radically new but always to seek perennial iterative modification, so that ‘it is never a process that begins from scratch: to design is always to redesign.’ It is for this reason – the practices of design as open-ended iterative modification, even as ‘anti-revolutionary’ – that I align contemporary machine learning models with design and not strictly with engineering.⁶⁴

Indeed, many contemporary deep learning practices such as ‘transfer learning’ definitively reject ‘handcrafting representations’, in favour of ‘greedy exploration’.⁶⁵ Every action is a modification of the residue that is already lodged within the layers of the model, it is never a complete overhaul or disruption. As Latour suggests, ‘to say everything has to be designed and redesigned, it will never be revolutionary’.⁶⁶ This foreclosure of something different, something revolutionary, is a crucial problem in machine learning political epistemes. As I have described, even where the practice of trialling a model is found to be in breach of law, or where racist data inputs are removed, still nothing revolutionary can emerge. For design can always modify and adjust and move the threshold, each adjustment another indelible mark, a marking out and a demarcation line. When Latour concludes that ‘designing is the antidote to founding, colonizing establishing’, I must disagree with him, for it is precisely colonising in ways that incorporate ever-increasing layers, extend to ever more domains of life, and dwell quietly in the violences of the modified weight. What new or alternative politics can possibly emerge when every potential pathway has already been narrowed to a mere parameter? It is to the implications and potentials for alternative political futures that I now turn in conclusion.

Design interruptions: Resisting machine learning worlds

In setting the themes that animate this Special Issue on disruption, Nicole Grove posed the question ‘what kind of worlds are in store for us as algorithms disrupt forms of organisation and advocacy for more equitable futures?’⁶⁷ I have sought to map out how machine learning actively incorporates the data from disrupted and fractured forms of organisation, and why it is that advocacy for alternative political futures becomes foreclosed in the logics of retroactivity and perennial trialling. I have suggested that a machine learning political episteme – one that eschews rules-based computational and political orders – is profiting from the undoing of postwar international and social institutions, from the deep neural networks powering the Vote Leave campaign to the so-called ‘digital transformation’ of the pandemic NHS. While, of course, I am not nostalgic for cybernetic worlds of rules-based computation and the liberal international order, nevertheless it is the case that notions of democratic life, human rights, and social ethics also grew amid such rules oriented orders. Where machine learning political orders are precisely profiting from the undoing of rights and collective public institutions, there are new challenges for the politics of resistance.

What happens to the space for resistance amid the power of the machine learning algorithm? What are the possibilities for reopening the futures that are condensed and foreclosed in the output of a deep neural network? Where machine learning algorithms are increasingly learning from the features of social scenes and the gathering of people in public space, is collective politics reduced to a being together that is merely the clustering of attributes? As Judith Butler has put the question, ‘what does it mean to act together when the conditions for acting together

⁶³Ibid., pp. 3–4.

⁶⁴The geographer Kevin Grove describes ‘resilience thinking’ as animated by what he calls a ‘will to design’, which works with the grain of uncertainty and experimental futures. Grove, *Resilience*.

⁶⁵Yoshua Bengio, ‘Deep learning of representations for unsupervised and transfer learning’, *Journal of Machine Learning Research*, 27 (2012), p. 28.

⁶⁶Latour, ‘A Cautious Prometheus?’, p. 8.

⁶⁷Nicole Grove, ‘Engineering Disruption’, Working Paper July 2020. See also Grove, this Special Issue.

are devastated or falling away?'.⁶⁸ Such questions are more urgent and acute because the threats to the rights to protest and freedoms to assemble are intensified by a machine learning order that absorbs the attributes of collective action. In her treatise on political assembly, Butler imagines that the 'gathering signifies in excess of what is said' and that 'popular assemblies form unexpectedly and dissolve', exercising a 'plural and performative right to appear'.⁶⁹ Yet, this plural and performative excess of the gathering of vulnerable bodies in public space is precisely under threat from the retroactive and trialling logics of the machine learning polity.

When the machine learning algorithm becomes the *mise-en-scène* of the public square, the means of arranging the scene and extracting the features, what political claim can be heard that is not already extracted and scored, and who can make it? The task for resistance, I suggest, is to interrupt the ordering of the political scene in order to ask how it might be otherwise. My emphasis on interruption consciously rejects the vocabularies of disruption that animate the force of disruptive technologies and 'push on the fracture until it breaks' tech industry cultures. To interrupt the scene is to resist its very condition of appearance, to locate the breaches in algorithmic arrangements and to show how they could be otherwise. As Walter Benjamin notes of Bertolt Brecht's device of 'interruption' in epic theatre, 'the truly important thing is to discover the conditions of life', where this discovery 'takes place through the interruption of happenings'.⁷⁰ To interrupt the scene of a machine learning political order would be to confront the plural branching pathways that could have yielded a different output and to amplify those branches as political decisions. In every arrangement of a machine learning model there are the traces of the rejected alternative. Brecht's device of interruption presents the observer with the traces of what could have been present, with the actor performing 'in such a way that the alternative emerges as clearly as possible', allowing 'other possibilities to be inferred' even while she 'represents one out of the possible variants'.⁷¹ In this way, the interruption of the scene works against the grain of the algorithm's reduction to one visible output, showing the contingency and multiplicity of the one out of many possible variants. Here lies a significant form of resistance; to amplify the branching points as moments where things could have been otherwise, where other possibilities could be inferred; and to refuse the reduction of political difficulty to one that is the output.

To resist being governed by a machine learning political order will necessitate naming the harms – beyond the conventions of privacy, data protection, and existing bodies of rights – of the foreclosure of alternative political futures. Though the machine learning political orders I have described close off political contestation and unheard claims, under the figure of the machine learning model there remains a teeming politics. When the solution precedes the political problem, the adjustment of parameters is also a real and violent modification of people's lives – as migrants, as benefit claimants, as people gathering in the city square. It is for this reason that the deep learning practice of modifying 'weights' in the model must be rendered heavier and more burdensome than the lightness of an adjustment implies. The weight in a machine learning model is not merely a technical weight on a connection in the neural net. It is the full burden and heaviness of a rejected visa application, a past facial biometric captured at a protest, a refused welfare claim, the extracted features of the refugee. In her compelling account of how colonial formations endure, Ann Laura Stoler foregrounds the 'enduring fissure' and the 'durable mark' of imperial duress.⁷² Stoler's affecting thought about 'duress' foregrounds the 'hardened, tenacious qualities of colonial effects' and 'endurance' in the 'capacity to "hold out" and "last", to endure as a countermand to "duress" and its damaging qualities'.⁷³ The weight of the machine learning algorithm could be freighted with the heaviness and endurance of Stoler's imperial duress.

⁶⁸ Butler, *Notes Toward a Performative Theory of Assembly*, p. 23.

⁶⁹ *Ibid.*, pp. 7, 11.

⁷⁰ Walter Benjamin, *Illuminations* (New York, NY: Schocken Books, 1999), p. 147.

⁷¹ Bertolt Brecht, *Brecht on Theatre* (London, UK: Methuen, 1986), p. 137.

⁷² Ann Laura Stoler, *Duress: Imperial Durabilities for Our Times* (Durham, NC: Duke University Press, 2016), p. 6.

⁷³ *Ibid.*, p. 7.

Each adjustment and modification of the model a squeezing and a tightening of the conditions of liveability of a political space, a community. Every indefinite trial a trial in the fullest sense of something that is borne by vulnerable bodies.

Acknowledgements. This article developed as part of a collective and collaborative forum led by Dr Nicole Grove. In July 2020 Nicole wrote to a group of us with a proposal to engage with the theme of ‘Engineering Disruption’, suggesting how ‘concepts of design and engineering have become themselves forms of political intervention rather than a means of political intervention.’ In the months that followed, Nicole led us in a series of extraordinary conversations and discussions. I am deeply grateful to Nicole, to Nisha Shah and Martin Coward at the *Review of International Studies*, and to Charmaine Chua and Neel Ahuja for the wonderful and enduring discussions we had in the group. Earlier versions of the article were presented at the Engineering Disruption workshop in November 2020, where Rocco Bellanova, Carola Westermeier, Debbie Lisle, Maja Zehfuss, and Martina Tazzioli so generously gave their insights; and at the Disruption by Design roundtable at the British International Studies Association conference, June 2021, where the discussion of error and failure with our audience was particularly generative. I am grateful to the three anonymous reviewers for *RIS*, whose comments have been so very helpful to me, and the Editors Martin Coward and Nisha Shah for their patience and generous depth of engagement. The research has received funding from the European Research Council (ERC) under Horizon 2020, Advanced Investigator Grant ERC-2019-ADG-883107-ALGOSOC and from the Independent Research Fund Denmark for the AI Reuse project.

Louise Amoore is Professor of Human Geography in the Department of Geography, Durham University, UK. She works on the politics of algorithms, the geopolitics of technology, biometric futures, and the ethics of machine learning systems. She is the author of *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others* (Duke University Press, 2020), and *The Politics of Possibility: Risk and Security Beyond Probability* (Duke University Press, 2013). Louise’s current research is funded by a five year ERC Advanced grant, ‘Algorithmic Societies’.