CONVERGENCE

Convergence: The International Journal of Research into New Media Technologies
2022, Vol. 0(0) 1–20

DOI: 10.1177/13548565221133248
journals.sagepub.com/home/con


Cynical technical practice: From AI to APIs

Sam Hind
University of Siegen, Germany

Tatjana Seitz
University of Siegen, Germany

Abstract

In this article, we examine how critical thinking, methods and design are used within the tech industry, using Philip Agre's notion of critical technical practice (CTP) to consider the rise of 'cynical' technical practice. Arguments by tech firms that their AI systems are ethical, contextual, situated or fair, as well as APIs that are privacy-compliant and offer greater user control, are now commonplace. Yet, these justifications routinely disguise the organisational, and economic, reasons for the development of technical systems and features. The article considers how different forms of 'technical critique' are used by technical practitioners such as software engineers, applying Agre's work on CTP, AI planning, grammars of action and empowerment to evaluate, and contextualise these justifications. As Agre understood, technical practitioners are not necessarily 'a-critical' or 'uncritical' in their approach to the design of technological systems or methods, but ordinarily compare the utility or performance of such according to a golden ethic: 'does it work?'. Drawing on Agre's studies of AI in the 1990s, the article considers how and what Agre considered to be the 'Cartesian soul' of AI research, on linguistic structuralism, and continues to frame much work within the wider tech industry today. Yet increasingly, as the article shows, 'narrow' and cynical forms of technical criticality are being used to legitimise, and publicise, corporate strategies of tech firms, whether through the development of AI systems by automotive start-ups such as Comma, or the management of relations with external developers through APIs, in the case of Facebook. Rather than judging the moral character of technical practitioners, however, the article offers an approach – via the work of Philip Agre – to examine how critical thinking is used, and often abused, within and beyond the tech industry.

Keywords

Critical technical practice, cynical technical practice, AI, APIs, grammars of action, empowerment, privacy, power

Corresponding author:
Sam Hind, SFB 1187 Media of Cooperation, University of Siegen, Herrengarten 3, Siegen 57072, Germany.
Email: sam.hind@uni-siegen.de

Introduction

Philip Agre first formulated the concept of 'critical technical practice' (CTP) in *Computation and Human Experience* (Agre, 1997a), in his efforts to 'reform' AI (Agre, 1997b) during his PhD at the MIT AI Lab. It was here that Agre worked on situational AI 'planners', building an artificial agent called Pengi with David Chapman (Agre and Chapman, 1990) that could play a basic computer game. Through Pengi, Agre and Chapman proposed a generative approach to AI, 'offering an infinity of dynamic possibilities rather than an infinity of structural combinations' (Agre and Chapman, 1990: 23) common to classical, rule-based approaches to AI. Pengi, thus, can be seen as one of the first software methods, and the study of the 'conditions of possibility that software establishes' (Fuller, 2008: 2).

Whilst Agre demanded methodological synthesis, with 'one foot planted in the craft work of design and the other foot planted in the reflexive work of critique' (Agre, 1997b: 155), he also offered thematic plurality. Four years after Agre and Chapman outlined their situational approach to AI, Agre published work on the more rudimentary technologies embedding themselves in contemporary workplaces, restructuring work-related activities according to so-called 'grammars of action' (Agre, 1994). Rather than considered as two discrete moments of 'early' Agre on AI, and 'later' Agre on grammars, we argue that these different strands of Agre's work can be brought together.

Fuelled by datafication (van Dijck, 2014), platformisation (Helmond, 2015) and infrastructuralisation (Plantin et al., 2018), whether big tech works with AI planners or API grammars these inventive solutions, we argue, are often framed as 'critical' technical interventions, offering 'better' or 'optimized' services for the good of a user community and wider society. These interventions rest on the belief that technical implementations, designs, or frameworks are built critically, folding in broader questions around ethics, bias and representation in the design, for instance, of AI object-recognition processes or API-reliant privacy protocols.

However, as we subsequently argue, such work is typically framed as either narrow or broader kinds of critical practice mentioned by Agre (1997b). In so doing, they constitute forms of technical criticality we refer to here as *cynical* technical practice. Rather than offering a kind of 'intellectual vagueness' (Jaton, 2021: 2) masking real intentions, technical criticality is cynically wielded by tech firms to accumulate, and consolidate, organisational power. Agre's formulation of a critical, technical practice, we argue here, can help evaluate how, and where, this discursive work happens within the tech industry.

The article explores two cases exhibiting different depths, kinds and orientations of this cynicism. At one end, automotive start-up Comma's driver assistance products, and at the other end, Facebook's Application Programming Interface (API), the Graph API and permissions architecture. In the former, we contend that technical criticality is cynically used to convince users of the 'empowerment' that a data-driven approach to AI affords them, whilst 'cherry-picking' performance metrics to sell the device itself. In the latter, we argue that criticality is cynically used to reorganise relationships with third parties including competitors, as well as to enforce and self-police regulatory requirements.

Agre's pioneering work on planners and grammars allows us to make sense of AI platforms that are framed as 'reflexive' or 'contextual' to APIs that are considered 'privacy-compliant' or 'sensitive' to user privacy. In other words, to evaluate how criticality is used by tech firms to demonstrate to users, and justify to wider social actors (governments, regulatory bodies, judicial systems, unions, activist groups), their apparent commitment to critical thinking, critical methods and critical design. In this, Agre was prescient with his work affording us greater conceptual and

methodological opportunities to explore how organisational power is accrued through the work of technical practitioners.

In the first section of the article, we discuss critical trends in the development of AI and APIs. In the subsequent section, we discuss the nature of technical critique and introduce the notion of cynical technical practice. In the next two sections, we provide a reading of Agre's work on AI planning and grammars of action, two technical mechanisms that enable this accrual of organisational power. Tying these back to the limits of reflective practice, we then offer an interpretation of what Agre meant by critique. In the final two sections, we provide evidence of this cynical criticality, with the cases of the automotive start-up, Comma and Facebook's Platform for Developers.

Critical AI to critical APIs

CTP has been taken up by a broad community, with the birth of various critical studies across the social sciences and the humanities, from critical data studies (Dalton et al., 2016; Iliadis and Russo, 2016) to data journalism (Bounegru and Gray, 2021) and digital art practice (YoHa, 2021), expressly building on Agre's call for 'critical methods' (Agre, 1997b: 132) in fields that typically work with, or design, technical systems. Here, the idea of critical technical practice that emerged from Agre's (and Chapman's) original work on situated AI has subsequently been folded back into contemporary thinking about algorithms and automation, as well as work on APIs and social media platforms.

Regarding AI, ideas around contextual, situated systems are increasingly common. Calls for 'ethical AI' (Blackman, 2020) often concern the operationalisation of 'fairness' (Veale and Binns 2017) and the identification of risks 'such as biased algorithms, privacy violations, and unexplainable outputs' (Blackman, 2020: n.p.), that are often either meant to soften hard-coded categories or processes (contextual) or offer expanded categories or more rigorous processes (situated). Other calls for 'rethinking' AI (Fuchs and Reichert, 2018), the 'radical democratisation' of AI (Verdegem, 2021: 12), or a 'decolonial' AI (Mohamed et al., 2020) have concerned issues around accessibility, opportunity and representation, whilst questioning the extent to which AI is for 'social good' (Verdegem, 2021: 14) as big tech firms have sometimes argued in the generalised turn to AI (Katzenbach, 2021). Indeed, in the case of Mohamed et al. (2020), co-authored by an employee of DeepMind, there is even an explicit call to shift 'towards a critical technical practice of AI' (Mohamed et al., 2020: 672).

Kate Crawford and Ryan Calo (2016: 313), likewise, suggest a 'social-systems' approach to AI meant to integrate social concerns at each stage of the design process, from conception to regulation. In this, they suggest that such an approach could 'ask whether the risks and rewards of the system are being applied evenly' (Crawford and Calo, 2016: 313), and 'consider the social and political history of the data' (Crawford and Calo, 2016: 313) used to train AI systems. Recent efforts to offer a 'new materialist AI' (McQuillan, 2021: 75) also critique the extent to which contemporary AI is an 'engine of injustice' (McQuillan, 2021: 75) in its drawing of boundaries and borders around (and through) people, instead offering a 'post-human' version (Schwartz, 2018) that recognises its agentic qualities, whilst limiting its descriptive power to reduce, represent and universalise.

Some of this work is interested in refining AI rules and plans, other aspects of it are concerned with the contextual intelligence of any AI system. In any case, calls for 'ethical AI' or similar tend to demand a widening, and deepening, of the social and cultural *knowledge* of AI such that it is able to account for difference, variety and specificity. In other words, such technical systems are indeed designed with 'critical methods' in mind.

There have also been efforts to establish critical interpretations of social media APIs. Benjamin Grosser's *Facebook Demetricator*, for instance, removes the quantitative dimension of using the platform (i.e. through 'like' metrics) in order to explore how they change the user experience (Grosser 2014). More recent projects by Grosser, including *Platform Sweet Talk* (2021), which explores the tricks used by platforms to stimulate engagement, and *Minus* (2021), a social media platform limiting users to 100 posts, are both aimed at critically appraising the user interactions that APIs – and specifically API grammars – organise and formalise (Grosser 2021a, 2021b). Jason Chao's 'AppTraffic' tool (Chao, 2020) likewise seeks to capture, inspect and playback the data sent between mobile apps and other parties (hence, app 'traffic'). In a more conceptual vein, Taina Bucher has considered how the Twitter API functions as a 'protocological object' (Bucher, 2013: 3) controlling both the exchange of information as well as wider social relations. Similarly, Eric Snodgrass and Winnie Soon (2019: n.p.) have argued that APIs in general function as 'infrastructural elements, facilitators and arbitrators', and as such, open up a particular space for critical interpretation and use.

However, contemporary efforts within the tech industry to reform both AI and APIs are still marginal, as the firing of Timnit Gebru, co-leader of Google's Ethical AI Team, suggests (Metz and Wakabayashi, 2020). Agre suggested in 1997 that no 'coherent "critical" school ha[d] arisen' (1997b: 154) to challenge the 'classical' school of AI which arguably still stands today, and critical attempts to reform how APIs function are largely limited to artistic critiques. These forms of critical technical practice (henceforth, critTP) can, we argue, be distinguished from what we call here *cynical* technical practice (cynTP). Our aim is not to label all classical AI and API cases as necessarily 'cynical', but to identify a dominant discourse within the tech industry that mobilises a narrow form of technical criticality. We will now talk about what this cynical form of technical practice involves.

Cynical *technical practice*

The term 'cynical' has two definitions. The first is to be 'contemptuously distrustful of human nature and motives' (Merriam-Webster, 2021: n.p.). The second is 'based on or reflecting a belief that human conduct is motivated primarily by self-interest' (Merriam-Webster, 2021: n.p.). In the second definition, to be cynical is to be singularly concerned with doing something, at the expense of other people and societal norms. For example, in a way, a politician might win votes by appealing to environmental concerns, despite not holding any 'green' values. Our definition of cynicism, however, is not primarily concerned with the moral integrity of individuals nor with the Ancient Greek philosophy of the 'cynics', but with an organisational capacity to claim critical thinking, methods or design for corporate self-interest, at the expense of some form of social good or justice.[1]

In contrast to critTP, then, cynTP might be seen as a kind of 'corporate' research (Langlois and Elmer, 2013), performed largely or only by firms. In Langlois and Elmer's reading of mass media research in the 1930s and 40s, critical research 'aimed to identify systemic power inequalities as well as formulate alternatives' (Langlois and Elmer, 2013: 6). In contrast, corporate research largely 'aimed to be descriptive', being both 'state- and commercially-driven' (Langlois and Elmer, 2013: 6). Such research was principally oriented towards consumers and consumer products, intended to offer methods to better understand consumer interests and, ultimately, design better products. In the case of social media platforms, this might be the difference between the critical work of Grosser (limiting, or 'de-optimizing', usability) and work by researchers funded by, or working for, big tech companies, such as the internal reports released in the 'Facebook Files' (Wall Street Journal, 2021).

In their call for a critical data studies, Dalton et al. (2016: 1) suggest that referring to work as critical risks 'bifurcating' critical theory and empirical work. In this, the gap between 'critical' research and corporate research is less obvious, even if they suggest that work in critical data studies must, ordinarily, 'contest the creation, commodification, analysis and application of data' (Dalton et al., 2016: 1), making space for the 'recursive dialog between the deeply theoretical and the robustly empiric' (Dalton et al., 2016: 1). This 'recursive dialog' seems to suggest that any gap between critical (theoretical) work and (implicitly) 'non-critical' work is better formulated as a set of non-opposite poles, emphasising the need for a dialectical approach to studying data. In this, theoretical work and empirical work should work in tandem. 'More specifically', as Agre (1997b: 15, authors' emphasis) sums up, 'the object of critical reflection is not computer programs as such but the *process* of technical work'.

However, we contend that cynTP is somewhat different in that it still offers a version of critical work. This is not simply that technical approaches or systems are referred to as 'situated', 'reflexive' or 'ethical', but that they are designed according to some epistemic logic. This kind of criticality was recognised by Agre himself, suggesting that the often 'only legitimate form of critical argument' offered by AI practitioners is that '"my system performs better than your system on problem X"' (Agre, 1997b: 150). As he explains, AI does not have 'ideas' in the sense philosophy might have, but 'technical practices' as well as 'loosely articulated intuitions' (Agre, 1997b: 150) about them.

Following this, Agre suggests that 'AI's construction of itself as a self-contained technical discipline' (Agre, 1997b: 150) is 'governed by practical-minded criteria of success and failure' (Agre, 1997b: 150) that, in contrast to what its proponents believe, 'is actually a powerful force for intellectual conservatism' (Agre, 1997b: 150). Here, Agre's idea of a critical technical practice, stimulated by what Malik and Malik (2022) call 'critical technical awakenings', is a response to this conservatism, which limits criticality to the comparative evaluation of the performance of technical systems.

Yet, Agre acknowledges that these people in AI regard these comparative activities as both legitimate *and* critical, forming their own 'clearly discursive construction' (Masís, 2014: 61). Whilst such work may indeed be narrow, and still concerned with technical performance, rather than wider ideas around knowledge, perception, independent agents or situations, AI practitioners nonetheless regard them as 'critical' in some sense.

As a technical discipline work in AI research, Agre argues, is organised around specific activities: 'to do AI is to prove theorems, write software, and build hardware whose purpose is to "solve" previously defined technical "problems"' (Agre, 1995a: 1). 'The whole test of these activities', Agre writes further, 'lies in "what works"' (Agre, 1995a: 1), and whether a technical resolution for 'problem X' can be found.

In this, Agre (1995a: n.p.) considers the 'Cartesian soul' of AI, from Descartes to Chomsky to the American psychologist Karl Lashley. Lashley in particular not only sketched the 'shape of a future problem' but also 'the principal strategy of a whole generation for solving it' (Agre, 1995a: n.p.), establishing that all actions could 'be understood on the model of language' (Agre, 1995a: n.p.). Work within AI, as Agre understood it, thus 'consisted in formalizing, elaborating, implementing, and testing those ideas' (Agre, 1995a: n.p.) formulated through this foundational linguistic structuralism.

Yet Agre wanted AI to be more than the testing of technical criteria, such that AI should provide, 'powerful means of forcing into the open the internal structure of a system of ideas and the internal tensions inherent in the project of getting those ideas to "work"' (Agre, 1995a: n.p.). In other words, that AI needed to be more critical. Whilst Agre (1995b; 1997b) mobilises this argument to propose a wider technical-discursive approach, it seems right to investigate what this rather narrow – yet, still

legitimate to most practitioners – critical dimension looks like in practice, in the development of AI and APIs.

To reiterate, Agre suggests that AI practitioners claim the work they do, and believe the evaluative practices they engage in, pass for criticality. In this, such work operationalises a narrow form of *technical criticality*, often referred to as 'technical solutionism' (Morozov, 2013) in corporate settings, that does not mobilise the same level, or kind, of critique that Agre (1997b) argues for. As Michael Dieter (2014: 218) suggests, the idea of 'problematization', taken from Foucault, concerned a 'repositioning of problems' that involved both discovering and *creating* them as articulative objects. In this, the 'critical' aspect of technical practice involves the 'suspension of "means" and "ends" across sociotechnical experience' (Dieter, 2014: 218), opening up the craft/critique aspect of critTP. In a narrow form of technical criticality in which evaluative practice pass for criticality in the eyes of practitioners, 'means' and 'ends' remain.

Focusing on the discursive dimension of the US tech industry, Fred Turner (2006) famously considered the intertwining of 'New Communalism' with engineering culture in Silicon Valley: where utopianism originally met technical solutionism. Yet, whilst the engineer's ethics of 'does it work?' (Turner, 2017: n.p.) serves the interests of other AI researchers within a self-contained framework ('narrow'), now corporate tech presents technically 'narrow' frameworks to legitimise managerial business policy ('cynical') (Levy, 2020).

In recent years, the disturbance of this narrow logic has left engineers at big tech firms confused when confronted with the perceived harm caused by otherwise 'elegantly' built systems (Turner, 2017). However, our interest in this narrowness is not to judge the moral character of technical practitioners, but to engage with this 'gap' that causes irritation. Following Agre, our interest is in examining the organisational power that constrains the possibility of critical thinking, methods, or design becoming part of the work process. In this, we understand criticality as a technical practitioner's capacity for individual decision-making in their ordinary working practice, routinely dependent on their role and position within an organisation.

Such constraints are evident in AI research settings as Agre (1997b) suggests, but also in domains in which AI is deployed, and APIs are developed. Here, technical criticality structures the relations between practitioners in the field (and domains) of AI and APIs, protecting the boundaries of technical work and solidifying organisational hierarchies. We refer to this as *cynical* technical practice (cynTP) in the way it cynically presents itself as critical, as evidenced in some of Agre's own work on the 'empowerment' of computerised workplaces (Agre, 1995b). CynTP allows us to evaluate whether technical practitioners are in control of the design decisions they make (Comma), or whether social media users are in control of their privacy and personal data (Facebook).

CynTP typically fits both definitions of cynical in that purveyors are both *distrustful* of 'alternative' theories that do not pass a means/ends, success/failure evaluation and cynically mobilise criticality for *self-interest*: to protect the specific form of technical critique engaged in, to justify the work being performed and to 'fix' (Katzenbach, 2021) social problems through technical means. Within the tech industry, this self-interest naturally manifests as different kinds of political, social and economic power. At its critical end, whistleblowers offer a puritanical defence of their former big tech employers (Milmo, 2021), whilst at its more cynical end, practitioners valiantly try to square an ethical circle.

As both of our cases intend to show, ideas around reflexivity, responsiveness, situatedness and inventiveness are just as common to cynTP. In this, resolving the 'Problem X' question in AI and API work is not always explainable through reference to a notion of acritical corporatism, as the Facebook Files story supports. Instead, it can be better understood through how a narrow, technically-oriented criticality ('means') entrenches organisational power ('ends'). In the following

we will explore two technical mechanisms discussed by Agre that operationalise this technical criticality: *planners* and *grammars*.

Planners (and decisions)

Agre and Chapman's (1990) work on AI planning centred around two kinds of computational models.[2] The first, known as the 'plan-as-program' perspective (Agre and Chapman, 1990: 17) designs an artificial agent to follow a pre-formatted plan in a '"mechanical" fashion' (Agre and Chapman, 1990: 18). In this, 'primitive actions and monitoring conditions' are carried out by the agent, but with 'little or no new reasoning about the activity' (Agre and Chapman, 1990: 18) at hand. Following such a model, the agent merely follows orders it is programmed to. In effect, the AI at play here is rather simplistic: a mindless body executing the plans of a disembodied planner.

Agre and Chapman detail the fundamental decision-making problems with this approach. Firstly, that it presents 'computationally intractable problems' (Agre and Chapman, 1990: 18) stemming from the need to formalise an exhaustive number of 'undecidabilities' (Agre and Chapman, 1990: 18) to which the agent must follow. Secondly, that it is incapable of working in any environment 'characterized by unpredictable events' including 'the actions of other agents' (Agre and Chapman, 1990: 18). Thirdly, that it requires an almost infinite regress of detailed plans and planning conditions. Then, fourthly, that it fails to provide a convincing connection between the written plan and the 'concrete situation' (Agre and Chapman, 1990: 20) the agent would typically encounter. As Masís (2014: 62) summarises: 'AI can go on well without a well-developed concept of the environment but only at the price of focusing on mere toy-problems, microworlds, and toy-tasks within such artificial environments'. In a world devoid of contingency, the plan-as-program agent is king.

In contrast, Agre and Chapman proposed a 'plan-as-communication' approach (Agre and Chapman, 1990: 24) that served as a prototype for critTP. In this, an agent does not 'mechanically execute' (Agre and Chapman, 1990: 17) a pre-written plan, but improvises based on the circumstances. One of the guiding principles is the *situatedness* of activities, in which '*all* activity…takes place in some specific, ongoing situation' (Agre and Chapman, 1990: 25, authors' emphasis). In this, a plan-as-communication view is designed to better negotiate unpredictable environments, or tackle 'ill structured problems' (Simon 1973; Ribes et al., 2019). Under a reflexive approach, this unpredictability becomes a resource in which 'contingencies … can occasion creative improvisation' (Agre and Chapman, 1990: 26).

Each model represents the world differently. In the plan-as-program model, the world exists external to the planner, what Masís (2014: 65) refers to as 'objective representation'. In this, Agre and Chapman's artificial agent (Pengi) would encounter other entities (such as two different bees) that would be definitively labelled (BEE-34, BEE-35). Under a 'deictic' model, however, the bees assume different labels based on where, and how, the agent encounters them in navigating the game space. For instance, BEE-34 might become '*the-bee-I-am-chasing'* (Agre and Chapman, 1990: 21, authors' emphasis). Put in decision-making terms, the model offers an 'ontogenetic' approach to the world, 'brought into being' (Kitchin and Dodge, 2007: 335) by Pengi's interactions, with 'routines and activities…realized "out there" in the world' (Masís 2014: 65).

This generative, ostensibly *critical*, approach to AI is meant to resolve the four problems. Here, 'undecidabilities' do not need to be formalised, as the plan itself is not hard-coded. Likewise, (other) unpredictable agents in contingent environments do not represent a problem, as the agent is able to render the 'chaos' intelligible (Norman, 1993). As a result, there is no hierarchy of subplans, which typically allow the (human) planners to live in a 'simple, abstract world' (Agre and Chapman, 1990: 19).

Instead, the agent is left to figure out things on its own, learning from its encounters, making 'fresh decisions' (Agre and Chapman, 1990: 23). The fourth problem disappears as a hard-coded abstraction of an environment is replaced with a 'situational-specific improvisations' (Agre and Chapman, 1990: 20) based on domain knowledge derived from the field. As Agre and Chapman (1990: 20) conclude, the generative approach views an agent as '*participating* in the flow of events' (Agre and Chapman, 1990: 20, authors' emphasis), rather than being fixated on solving problems.

In conclusion, this approach to AI planning can be seen as embryonic of critical AI, embodying the kind of sensitivity and contextuality desired from, but not always present in, contemporary AI. The aim of the model was not simply to prove another AI was technically possible, but that it could provoke critical debate, and discursive shifts, in the field itself. As Agre (1997a: 23) wrote: '[a]dvances in the critical self-awareness of technical practice are intellectual contributions in their own right, and … necessary conditions for the progress of technical work itself'. Having likely failed technical tests (Pengi was only ever developed under laboratory conditions), the approach nonetheless made a substantial epistemic contribution to AI research (Dreyfus, 2007).

Grammars (and accountability)

Agre also considered how computational systems were fundamentally changing human activities and forms of social organisation. To elucidate these rearrangements, Agre drew on Suchman (1992) and how workers at a regional airport 'accounted' for the work they do in managing aircrafts, passengers and baggage. As Agre suggests, this accountability 'is not just a formal relationship or an outside force, but a practical process of exhibiting reality' (Agre, 1995b: 182), in which workers are engaged in the 'process of representing the[ir] work' (1995b: 182), such that this accountability becomes work in itself. As this 'new style of work is heavily "staged"' (Agre, 1995b: 182), Agre extends Suchman's work through a historical analysis of the design of technology that make such work representable to the computer. Here, '[a]lthough many technologies are involved, distributed computing technologies play a crucial role in creating, storing, accumulating, manipulating, and transmitting … representations' (Agre, 1995b: 182). In other words, that computers actively, and continuously, shape practices of accountability.

Agre (1994; 1995b) was thus interested in how computers represented human activities through 'grammars' derived from vernacular language. For example, in how the grammar of restaurant activities are derived from the professional vocabulary used by waiters, chefs and managers, from 'orders' and 'change' to 'tabs' and 'tips' (Agre, 1995b: 183). As they represented workplace activities in particular ways, Agre (1994: 109) called them 'grammars of action', constituting a 'capture model' (Agre, 1994: 101) increasingly prevalent in multiple industries. His interest, thus, was the impact of such grammars on work *itself* and how workers made themselves accountable to the model.

Agre's work on grammars of action has been widely applied. Florian Sprenger (2018) uses Agre (1994) to explore how the motivations of individuals can be inferred through the data collection capacities of mobile devices. Van Doorn and Badger (2020) use Agre (1994) to understand food delivery as a series of connected activities (job selection, food pick-up, delivery etc.) subjected to a 'process of platform-mediated capture' (Van Doorn and Badger, 2020: 5) by app-based delivery services such as Deliveroo or Uber Eats. Meanwhile, Gekker and Hind (2020) identify the limits of Agre's capture model in the world of autonomous driving, and Gerlitz et al. (2019) consider how social media apps capture third-party activity.

In the capture model, Agre aims to incorporate the formal mathematical language, more representative of the software engineer, within the social realm of situational language. In this, a

grammar becomes a hybrid entity, both enabling computational functionality (through mathematical formalisation) whilst bearing social utility (through incorporation of social terminology). Computational grammars present themselves such that technical schemata appear identical to social schemas. In such cases, however, computers do not simply 'collect' data without disturbing the phenomena or activities that create it. Computer systems both capture work processes in a formalised manner and 're-inject' (Agre, 1995b: 184) it with operational value. In so doing, the computational representation, or grammar, 'becomes a resource in the activity itself' (Agre, 1995b: 183). As a result, this capture mechanism is 'never purely technical but always *sociotechnical* in nature' (Agre, 1994: 112, authors' emphasis), offering different kinds of operating power in the process.

As Agre (1995b) explored, grammars and capture mechanisms became a critical feature of new working practices in the 1990s, loosely built around the idea of 'empowerment'. In a business context, empowerment 'refers to a process by which employees are freed of bureaucratic constraints and given control of their work in order to make decisions and reorganize their local-work processes in accord with their own judgement' (Agre, 1995b: 170). A key facilitator of these processes was the desktop computer. Measurement, in the context of this regime, is the process by which the (work) activities of the 'empowered' employee are captured and fed back into the modulation, and management, of these activities. As Agre (1995b: 176) contended, whilst these two processes of empowerment and measurement were well-known within business, they were 'rarely identified as a single, coherent system'. Agre's synthesis codified the relationship between them. In other words, to not only contest the claim that empowerment is the freedom to make decisions, but to articulate how distributed decision-making is enabled by 'simultaneously centralizing control through measurement' (Agre, 1995b: 179). In conclusion, this approach to understanding how computational systems 'capture' and shape user activities, and how users ('lay' users and professional practitioners) are 'empowered' to use them, can be seen as the basis for critical API work that either refines, or challenges, the operation of contemporary platforms.

(More Than) The Reflective Practitioner

Whilst we do not provide a full account of Agre's use of critique, the two concepts of planners and grammars reveal that critTP is not a state to be achieved but rather a mode for technical practitioners to act in. Fuller (2008: 2), for instance, proposes criticism as a process that enables 'politics, society, and systems of thought and aesthetics' to 'renew themselves'. Donald Schön's (1983) influential book *The Reflective Practitioner,* cited by Agre (1995a), considers how reflection is practiced in various technical disciplines, from engineering to architecture. The subtle difference between reflexivity and criticality marks the difference between the 'end' of inscribing politics into technology and the open 'means' by which systems of thought can be mobilised.

Whilst reflexivity is integral to critTP, it is easy to fall into what Katy Waldman (2020: n.p.) calls a 'reflexivity trap', in which self-reflexivity is cast as a goal in itself. The peril of a reflexivity trap occurs when reflexivity becomes the purpose of criticality. Agre's efforts 'trying to reform AI' (Agre, 1997b) were not concerned with 'fixing' AI from a technical perspective nor about achieving some level of individualised 'self-reflexivity', but rather to enable practitioners to 'employ the tools of critical inquiry to engage in a richer and more animated conversation with the world' (Agre 1995a: n.p.). Here, Agre's critique concerned the impoverishment of the discourse within AI research where internal 'ways of thinking' had become prioritised (such as linguistic structuralism) almost entirely over 'external stimuli and their effects' Agre (1995: n.p.). AI research might have offered self-reflexivity, but at the expense of an open discourse.

The aim of Agre's AI research was not to present a 'workable' AI system. As 'a conversation with the world' critTP concerns how technology disrupts or disturbs social norms, ordinary work routines, and political capacities. Tracing the history of ideas across technical discourses is not intended to 'debunk' ahistorical technical meaning-making, but to 'prevent the passage to formalism' (Agre, 1995a: n.p.) that invariably shears technical work from their cultural underpinnings. Critique in critTP, then, is as much about the 'clarification' (Agre, 1995a) of social and institutional forms of relations as it is about 'renewing' systems of thought, or 'reforming' technology.

In the following, we apply Agre's work on planners and grammars in two cases, through the lens of critTP and cynTP. Firstly, on how AI planning is employed in the design of driver assistance software by automotive start-up, Comma. Secondly, on how API grammars and a permissions architecture are designed by Facebook. We argue that both provide evidence of cynical technical practice (cynTP), to accumulate, and consolidate, different examples of organisational power: from framing an 'end-to-end' (E2E) machine-learning approach as 'empowering' for drivers, and cherry-picking performance metrics (Comma), to implementing an app review process as a clandestine economic mechanism, and using a 'permissions' system as a self-regulatory enforcement tool (Facebook).

Case study #1: Comma three

Comma is an automotive start-up based in San Diego, California. Since 2015 it has been developing an open-source, advanced driver assistance system (ADAS), intended as an entirely new 'device class' within consumer electronics (Hotz, 2021, n.p.). Their Comma Three device is designed to plug into existing vehicle models (from Toyota to Honda), via on-board diagnostics (OBD) systems standardised in the automotive industry since the 1990s (Forelle, 2022). Similar in appearance to a dashcam, the Comma Three is powered by visual, and sensor, data captured by existing users. This data is then used to train the Comma AI itself.

At the first 'COMMA_CON' event hosted in July 2021, akin to Facebook's F8 or Google's I/O, company executives made their case for a generative, data-driven approach to AI. In so doing, they offered a form of cynTP. Firstly, Comma frame their E2E approach to AI as empowering for, and responsive to, Comma users/drivers whilst eliding its importance to their business model. Secondly, Comma 'cherry-pick' performance metrics to frame their Comma Three device as better than rival products. Together, these cynical applications of technical criticality are designed to further the organisational power of Comma within the consumer electronics, and automotive, industries.

E2E as empowerment. At COMMA_CON, Harold Schäfer (Comma Chief Technology Officer, CTO) provided justification for their approach to developing an ADAS (Schäfer, 2021). In Comma's case, AI planning is still an integral part of the operation, but it is not what Schäfer (2021: n.p.) refers to as 'classical planning' involving a near-infinite number of domain-specific rules and limitations such as 'stay in lane, don't hit cars, don't hit people, stop at traffic lights' (Schäfer, 2021: n.p.). This classical approach is also re-iterated by George Hotz (Comma President) when calling rival Waymo (2021) programmers 'coneheads' (Hotz, 2021: n.p.), referring to the continual need to write rules to account for miscellaneous road objects such as traffic cones. In such an approach, programmers are committed to near-infinitely expanding their 'prop stash' of peculiar objects liable to create problems for any autonomous vehicle (Madrigal, 2017; Hind, 2019). In this, Comma distinguish themselves from classical, rule-based AI projects.

Instead, Comma offers an E2E approach (Comma, 2021). This alternative approach does not require a perception phase in which its algorithmic system (known as 'openpilot') attempts to

classify objects from captured data. Instead, it directly feeds ever-increasing volumes of driving data derived from existing Comma devices into a planner. From the data it ingests, and with the help of multiple neural network services, openpilot predicts the path it would take. In such an E2E approach, there is no hard-coding of driving rules, no hand labelled ground truth process, nor any reading of lane lines or road signs. Instead, it offers a 'user-first' data-driven approach.

Schäfer effectively restates Agre and Chapman's (1990) problems with what they call the plan-as-program model: that '[p]erception models generally require hand labelled ground truth', that '[p]lanners are hand-coded with unbounded complexity', that the possible total list of information required by planners is 'endless', and that such a planner system is both '[d]ifficult to build' but also 'impossible to maintain' (Schäfer, 2021: n.p.). In this, Schäfer frames Comma as *critical*, responsive to the world it moves through, in which its openpilot AI not only 'leads a life' as Agre and Chapman (1990: 20) suggested, but also allows users to retain their control over, enjoyment of, driving.

Yet, in emphasising the value of this to users, Schäfer and Hotz both omit a significant detail: that free, easily accessible, automatically captured, 'E2E' data is integral to Comma's business model. Without Comma Three devices being used by a wide variety of drivers, driving different cars, in different locations, the openpilot AI achieves none of its situational powers. Due to the cost constraints of building AI systems (Luitse and Denkena, 2021), access to free data to train machine-learning models is invaluable for a start-up like Comma, otherwise reliant upon dedicated data collection fleets (like Waymo).

In offering a strictly technical framing of Comma's approach to AI, Schäfer and Hotz intend to convey Comma's commitment to critical thinking, methods and design. Yet this commitment is not done altruistically, as if Comma's sole, over-riding interest was to empower drivers. Thus, whilst Comma appears to offer a form of critTP, it does so by eliding the commercial case for doing so. As such, Comma depends on cynicality: offering continued agency to drivers, whilst selling the benefits of user- and data-driven automation back to them in form of the Comma Three.

'Cherry-picking' metrics. George Hotz has also suggested that the company will be 'level two, forever' (Hotz, 2021: n.p.), referring to the standard levels (1–5) used in the development of autonomous vehicles. In such a case, any Comma device will only ever be capable of performing specific kinds of longitudinal (forward collision warning) and latitudinal (lane following) tasks, rather than offering full (level 5) autonomy. Narrowing the conditions for AI success here allows Comma to represent their planning approach as viable, especially when compared with other projects.

In the first sense, this conditional success can be seen as cynical because of the way the 'Problem X' formulation is misused. In the comparisons made by Schäfer and Hotz with Waymo, the 'problem' is stated as two closely-related, but nonetheless, *different* problems of driver assistance (Comma) or full autonomy (Waymo). This imperfect, asymmetrical comparison undermines the golden rule of engineering ethics ('does it work?'), comparing two systems that not only use two different means (rule-based and data-based methods), but also result in two different *ends* (level 2 and level 5 automation).

Even when a fair comparison is made, certain metrics are cherry-picked to highlight partial operational successes (latitudinal control), whilst ignoring other failures (longitudinal control). Referring to a Consumer Reports (2020) test on competing ADAS, Hotz (2021) mentions how Comma beat all 18 competitors, including Tesla's Autopilot system. Yet, whilst it scored highly (9/10) on latitudinal tasks (lane keeping assistance), it performed poorly (3/10) on longitudinal tasks (adaptive cruise control). Whilst selectively choosing features that promote a company product over

a rival's is nothing new, it once again undermines the ethic at the heart of engineering work, as Agre (1995a; 1997b) considers, in favour of a boosterism important to Comma's business model.

These cynical comparisons are meant to persuade customers of the superiority of the Comma product. In doing so, however, they invalidate the terms of a narrow, technical criticality. Firstly, through an asymmetrical comparison with Waymo, in which Comma's data-based approach to driver *assistance* is deemed superior to Waymo's rule-based approach to full, vehicular *autonomy*. Secondly, that through a symmetrical comparison with rival ADAS, mention of poor performance on a key metric (longitudinal capability) is ignored. In these cases, Comma offers a form of cynTP: improperly, and deceptively, making technological comparisons to accumulate power in a nascent automotive market.

Case study #2: Facebook graph API

Facebook's Platform for Developers ('Platform') enables developers to interact with Facebook's data servers. The 'primary way' (Facebook, 2020: n.p.) to use these software products is through the Graph API. The Graph API (GAPI) is a highly formalised computational representation of grammars directly derived from user activities with, and on, Facebook. Such grammars consist of several entities: 'endpoints' which provide granular information about 'nodes' for objects, 'edges' for the relationships between objects and additional parameters such as permissions (Facebook, 2021). Examples of the formalised grammar of a photo-like-activity include 'id', 'gender' and 'user_friends' for the individual actor and 'user object', 'created_time' and 'location' for the image or 'picture object', itself. More than 100 data instances can be captured when someone likes a photo on Facebook. These flexible grammars are continuously updated, capturing user activity in real-time and 're-injecting' them into the frontend, offering users 'novel' ways to interact with Facebook.

Facebook's technological infrastructure provides the documentation of Facebook's internal organisational behaviour (Wu and Taneja, 2020), offering the potential for a political economic analysis (Agre, 1995; Helmond et al., 2017; Nieborg and Helmond, 2018; Plantin et al., 2018; van der Vlist et al., 2022). The second version of the new GAPI was designed to make user and developer activities *accountable* in economic terms. This offers a 'moment…of transition' (Agre, 1995b: 190) when the routinised and habitual is disassociated and reorganised, rendering the passage to renewed formalisation observable. Here the 'Facebook Leaks' documents (Campbell, 2018) can be used to make sense of the historic development, and utility, of so-called 'action grammars'.[3] Two elements of GAPI's action grammars, 'App Review' and 'Permissions' illustrate the shifts, in which *economic* decisions are cynically presented as only *technical* changes.

App review as economic mechanism. Formally, the 'App Review' refers to a procedure in which Facebook checks whether third-party developers fulfil certain criteria before they can access user data. Instead of a technical 'problem X', here, Facebook poses 'user experience' as the problem to which the data-based solutions are provided. The leaked emails show, however, the review was designed to offer an economic assessment of the value third-party developers added to Facebook 'at review time, before the app launches'. Here, user experience is only one factor for the estimation of value added.[4]

When Facebook launched the new App Review, it still lacked a solid Platform business model that could transform the Platform infrastructure into an asset and by extension into capital (Doganova and Muniesa, 2015). Internally known as the 'review tool', Facebook wanted to determine whether an app increased the company's value. If it did not, the app was classified as a 'data leach' and its requests 'rejected'.[5] Seeking value creation, Facebook introduced App Review as a

socially valuable privacy control, whilst constituting a key function for mobilising worth in the third-party ecosystem.

Up until this point, Facebook had few mechanisms they could call on to evaluate third-party app development. To fill that gap, Facebook staged the redesign of the 'Login' tool strategically in terms of empowerment that would give 'people more control over the information they share with apps' (Facebook, 2014). For Facebook, privacy had the potential to replace contractual obligation with 'social obligations' (Fourcade and Kluttz, 2020: 3). The trade press welcomed the discontinuation of access to friends' data, framing the manoeuvre as motivated by privacy concerns (Constantine, 2014).

The cynicality of these critical technological changes is further evidenced by a scene from Allison Hendrix's (Head of Facebook Policy) deposition.[6] When asked whether Mike Vernal (Vice President of Product and Engineering) used the word 'privacy' in an email thread discussing the design of 'privacy tools and controls', such as the App Review, Hendrix responded that he was 'talking about the Platform business model'.[7] When pressed about bad developer behaviour, inappropriate app content or user spamming (reasons Facebook used to justify the depreciation of access to friends' data) she answered similarly.[8] Here, user privacy appears as a convenient veil for changes to Facebook's revenue model.

Following Agre (1995b), here user 'empowerment' distracts from underlying, economic, judgements. Under the pretext of data protection, App Review installed a 'checkpoint' into the third-party ecosystem, disguising its concomitant economic value to the company. By inserting App Review into the developer's login flow, any app or service that wanted to access user data via GAPI was effectively required to implement the 'Login with Facebook' feature. In so doing, it fundamentally redesigned developer-platform relations. Cynically implemented, the promise to protect user data further consolidated Facebook's 'infrastructural power' (Munn, 2020) over both developers and users.

Permissions as regulation enforcement. Since 2012, the US Federal Trade Commission (FTC) has investigated Facebook's failure to protect user privacy (FTC, 2019). One measure central to the investigations has been the function of so-called 'granular data permission' (GDP) interfaces. On the Platform permissions, transform data points are used to formalise an activity into articulations that represent access proxies to specific activities. Developers encounter permissions when implementing the Login feature, to users they are presented as GDP displays during the login flow into external services and apps.

Multiple FTC rulings have considered how non-public user information is shared with external developers (FTC, 2011, 2012, 2020). To protect user privacy, the court rulings assign GDP displays a special role demanding Facebook 'obtain the user's affirmative express consent' (FTC, 2011: 5). The FTC requested Facebook inform users about how their data is shared with the ecosystem of external developers, whilst simultaneously ruling that Facebook had the 'right to inspect third parties' (Van Loo, 2020: 467).

Administrative permits, distinct from contracts, protocols and other related legal and administrative tools, are ubiquitous in modern society (Biber and Ruhl, 2014). In effectively allowing Facebook to issue administrative permits, this case evidences how 'policymakers have begun relying on third-party enforcement by the real gatekeepers of the economy: the firms who control access to core product markets' (Van Loo, 2020: 47).

Van Loo (2020: 471) shows that in 2012, Facebook was required to hire a third-party auditor who would evaluate Facebook's compliance (FTC, 2012). By 2020, Facebook was allowed to regulate and enforce sanctions upon third-party developers independently (FTC, 2020). Unlike an external

auditor, Facebook is a critical business partner of those developers that they are meant to regulate, enabling it to act as an extended regulatory body of its own ecosystem.

The FTC delegated to Facebook both the design and implementation of enforcement rules and sanctions, and compliance monitoring of third-party developers, to be performed through 'ongoing manual reviews and automated scans … or other technical and operational testing…' (FTC, 2020: 9). A democratic body that controls Facebook does not, of course, exist, meaning that tools developed by Facebook itself – namely, Login and App Review – were being used to ensure third-party compliance.

Policymakers in a range of industries from energy to drug administration, Van Loo critiques, increasingly rely for such third-party enforcement despite such firms being able to 'control access to core product markets' (Van Loo, 2020: 471, cf. Helmond 2015) through techno-regulatory arrangements. Facebook is thus a prime example of what Van Loo (2020: 471) refers to as an 'enforcer-firm': firms that enjoy officially legitimated authority to enforce sanctions on competitors whilst avoiding democratic mechanisms that curb their power. The permit power exercised by enforcer-firms provides the basis for social authority needed for the hegemonic sway of what Stuart Hall called a 'class alliance on behalf of capital' (Hall, 2019: 165).

Forging alliances involves a decade-long strategy designed to reinterpret socio-cultural issues *technologically*. For instance, Facebook offers technical solutions (such as App Review) to concerns over consumer data privacy. Permissions, understood as regulatory mechanisms in the FTC rulings, are imbued by Facebook with technical meaning around permissions, single sign-on (SSO) schemes for login, and GDP displays.

In routinely dressing up economic decisions as critical technical changes, largely around issues of user privacy, Facebook offers further examples of cynTP. On this level, politico-technical grammars enforce both the economic and hegemonic power of Facebook, offering new methods to control its dominant position in digital life.

Conclusion

In this article, we have examined the work of Philip Agre to understand a specific kind of technical work performed in the development of AI and APIs, that we have called here *cynical* technical practice (cynTP). The term identifies how technical criticality is used to advance the economic power of big tech firms: often used to justify technical 'fixes' to social problems, as well as constituting a discursive approach to the design, implementation and public presentation, of technical systems. Whilst Agre (1997b) talked at length about such an approach within the world of AI research, it was largely used to justify the development of a more critical approach, taken from his engagement with various strands of critical theory both at the margins of AI work in the 1990s, including sociologists such as Harold Garfinkel and Lucy Suchman, and continental thinkers from Foucault to Habermas.

However, Agre (1997b) acknowledged that this narrow technical approach did not lack a form of criticality in its operation, only that it was too concerned with 'work-ability', such that the only theories that could be deemed valid were those that could be successfully applied and 'made to work' in a technical sense. Such an approach has many active protagonists both then and now: practitioners who claim the work they do, and believe the evaluative practices they engage in, pass for criticality.

In this article, we have discussed two case studies where technical practice becomes cynical: in an automotive start-up called Comma, and at Facebook. These cases, we argue, constitute good examples of AI and API work, in which technical criticality is used to further corporate self-interest,

rather than social good or justice. In the first case, Comma's frame their E2E approach to AI as 'empowering' for users, eliding how integral it also is to their business model. Likewise, that Comma executives cynically 'cherry-pick' performance metrics, invalidating a golden engineering rule in the process: 'does it work?'

In the second case, we argue that Facebook mobilises a technical criticality to disguise their work to economically value user/developer activity, developing a permissions system framed around the notion of 'user privacy' and 'data protection', whilst using their power as a big tech firm to enforce relations with third parties. In this case, this cynicality is directed towards developers, users and civil society.

Our aim for this article was three-fold. Firstly, to synthesise Agre's earlier work on AI and later work on grammars of action, in order to make sense of the co-development of AI and APIs. Agre's work does not only remain hugely relevant to contemporary digital society, but in considering AI work and grammars, he was also able to establish a foundation for comparative analyses of the political 'valence' of technical work today. His emphasis, for instance, on situated planning approaches in AI, or the transformative effects of technical capture processes on human activities, offer a template for contemporary accounts of machine vision, the significance of human agency in AI, the role of APIs in governing external relationships, and how permissions architectures consolidate the hegemony of big tech firms.

Secondly, to evaluate the role of critical thinking, methods and design in the tech industry, as first examined by Agre (1995a, 1997a; 1997b) in AI research. In this, our desire was not to evaluate the 'morality' of technical work, but to use the notion of criticality as a tool with which to interrogate the organisational conditions under which technical practitioners justify, and frame, their work. What we have found in both cases, then, is different ways in which technical criticality is used cynically: to compare firms and rivals (in the case of Comma), to appease legislators (in the case of Facebook), or more generally to accumulate capital, increase product engagement or capture new users. These 'economies of virtue' (Phan et al., 2022: 121) deserve further investigation.

Thirdly, to advance understanding of the various concepts Agre has developed in his work: from situated planning to 'regimes' of empowerment. In each case, Agre has either developed rather specific meanings for general terms (grammars) or crafted genealogies of them (empowerment). Using these through the lens of both critical and *cynical* technical practice, we have considered the 'multivalence' (Gerlitz, 2016) of such concepts, their discursive deployment and epistemological value, within certain technical settings. Rather than providing a 'how-to' for critTP, our aim was to present it as a mode of thinking about the way '"technical" and "human" affairs [are] bound together within a dynamic tension' (Agre, 1995b: 190). In this, critTP offers a connective interpretation of what, at first glance, seems to be separate cultural, political, technical and economic phenomena. Agre's work on critTP, therefore, is critical to how we make sense of our saturated digital worlds, and how critical thinking is used, and often abused, within and beyond the tech industry.

Acknowledgements

The authors would like to thank the co-editors of the special issue, fellow participants in a corresponding author workshop, and two anonymous reviewers. Any errors, of course, remain our own.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: This work was supported by the Deutsche Forschungsgemeinschaft (262513311 - SFB 1187).

ORCID iDs

Sam Hind https://orcid.org/0000-0001-8347-3695
Tatjana Seitz https://orcid.org/0000-0002-7147-1278

Notes

1. Thanks to Michael Dieter for helping the authors clarify this point.
2. These models go by various names, which we use synonymously throughout this article. The plan-as-program model is also referred to as a 'rule-based' approach, or as 'classical' planning. The plan-as-communication model is also referred to as a 'generative' approach, a 'deictic' model or as 'situated' planning. In either, we use 'model', 'approach' and 'planning' interchangeably.
3. The Facebook Leaks documents contain internal communication between Facebook's senior management in which they discuss the major redesign of the Graph API.
4. Facebook Leaks, Exhibit 43, draft by Douglas Purdy (Director of Engineering) at the time.
5. Exhibit 43.
6. In a case against the former third-party developer firm Six4Three, in 2017.
7. Facebook Leaks, Exhibit 45, 2012; Subject: RE: [Open Graph PMs + EMs] Uploaded 2012_10_26 Platform data model v5.pptx.
8. Facebook Leaks, Exhibit 4; 1-617-542-0039, Deposition of PMQ of Facebook, Inc. Allison Hendrix, June 21, 2017.

References

Agre P (1994) Surveillance and capture: Two models of privacy. *The Information Society: An International Journal* 10 (2): 101–127.

Agre P (1995a) The soul gained and lost: Artificial intelligence as a philosophical project. Available: https://pages.gseis.ucla.edu/faculty/agre/shr.html (accessed 10 May 2022).

Agre P (1995b) From high tech to human tech: Empowerment, measurement, and social studies of computing. *Computer Supported Cooperative Work (CSCW)* 3 (1): 167–195.

Agre P (1997a) *Computation and Human Experience*. Cambridge: Cambridge University Press.

Agre P (1997b) Toward a critical technical practice: Lessons learned in trying to reform AI. In Bowker GC, Leigh Star S, Turner W, et al. (eds) *Bridging the Great Divide: Social Science, Technical Systems, and Cooperative Work*. Mahwah, NJ: Lawrence Erlbaum Associates, 131–158.

Agre P and Chapman D (1990) What are plans for? *Robotics and Autonomous Systems* 6(1–2): 17–34.

Biber E and Ruhl JB (2014) The permit power revisited: The theory and practice of regulatory permits in the administrative state. *Duke Law Journal* 64(2): 133–234.

Blackman R (2020) A practical guide to building ethical AI. Available: https://hbr.org/2020/10/a-practical-guide-to-building-ethical-ai (accessed 10 May 2022).

Bounegru L and Gray J (2021) (eds) *The Data Journalism Handbook: Towards a Critical Data Practice*. Amsterdam: Amsterdam University Press.

Bucher T (2013) Objects of intense feeling: the case of the Twitter API. *Computational Culture* 3: 1–28.

Campbell D (2018) Facebook leaks. Available: https://www.duncancampbell.org/facebookleaks (accessed 10 May 2022).

Chao J (2020) AppTraffic. Available: https://apptraffic.phil.uni-siegen.de/#/ (accessed 10 May 2022).
Comma (2021) *How Openpilot Works in 2021*. Available: https://blog.comma.ai/openpilot-in-2021/ (accessed 10 May 2022).
Constantine J (2014) Facebook is shutting down its API for giving your friends' data to apps. Available: https://techcrunch.com/2015/04/28/facebook-api-shut-down/ (accessed 10 May 2022).
Consumer Reports (2020) Active Driver Assistance Systems: Test results and design recommendations. Available: https://data.consumerreports.org/wp-content/uploads/2020/11/consumer-reports-active-driving-assistance-systems-november-16-2020.pdf (accessed 10 May 2022).
Crawford K and Calo R (2016) There is a blind spot in AI research. *Nature* 538: 311–313.
Dalton CM, Taylor L and Thatcher J (2016) Critical Data Studies: A dialog on data and space. *Big Data & Society* 3(1): 1–9.
Dieter M (2014) The virtues of critical technical practice. *Differences: A Journal of Feminist Cultural Studies* 25(1): 216–230.
Doganova L and Muniesa F (2015) Capitalization devices: business models and the renewal of markets. In Kornberger M, Justesen L, Mouritsen J, et al. (eds) *Making Things Valuable*. Oxford: Oxford University Press, 109–125.
Dreyfus HL (2007) Why Heideggarian AI failed and how fixing it would require making it more Heideggerian. *Philosophical Psychology* 20 (2): 247–268.
Facebook (2014) *The New Facebook Login and 2.0*. Available: https://developers.facebook.com/blog/post/2014/04/30/the-new-facebook-login/ (accessed 10 May 2022).
Facebook (2020) Facebook for developers, graph API. Available: https://developers.facebook.com/docs/graph-api/ (accessed 10 May 2022).
Facebook (2021) Graph API. Available: https://developers.facebook.com/docs/graph-api/overview (accessed 10 May 2022).
Federal Trade Commission (2011) Facebook agreement. Available: https://www.ftc.gov/sites/default/files/documents/cases/2011/11/111129facebookagree.pdf (accessed 10 May 2022).
Federal Trade Commission (2012) *Facebook Docket No. C-4365*. Available: https://www.ftc.gov/sites/default/files/documents/cases/2012/08/120810facebookdo.pdf (accessed 10 May 2022).
Federal Trade Commission (2019) FTC imposes $5 billion penalty and sweeping new privacy restrictions on Facebook. Available: https://www.ftc.gov/news-events/press-releases/2019/07/ftc-imposes-5-billion-penalty-sweeping-new-privacy-restrictions (accessed 10 May 2022).
Federal Trade Commission (2020) Facebook modifying order. Available: https://www.ftc.gov/system/files/documents/cases/c4365facebookmodifyingorder.pdf (accessed 10 May 2022).
Forelle MC (2022) The material consequences of 'chipification': The case of software-embedded cars. *Big Data & Society* 9(1): 1–12.
Fourcade M and Kluttz DN (2020) A Maussian bargain: accumulation by gift in the digital economy. *Big Data & Society* 7(1): 1–16.
Fuchs M and Reichert R (2018) Introduction: rethinking AI. Neural networks, biometrics and the new artificial intelligence. *Digital Culture & Society* 4(1): 5–13.
Fuller M (2008) *Software Studies: A Lexicon*. Cambridge, MA: MIT Press.
Gekker A and Hind S (2020) Infrastructural surveillance. *New Media & Society* 22(8): 1414–1436.
Gerlitz C (2016) What counts? Reflections on the multivalence of social media data. *Digital Culture & Society* 2(2): 19–38.
Gerlitz C, Helmond A, van der Vlist FN, et al. (2019) Regramming the platform: infrastructural relations between apps and social media. *Computational Culture* 7: 1–42.
Grosser B (2014) What do metrics wants? How quantification prescribes social interaction on Facebook. *Computational Culture* 4: 1–39.

Grosser B (2021a) Platform sweet talk. Available: https://bengrosser.com/projects/platform-sweet-talk/ (accessed 10 May 2022).

Grosser B (2021b) Minus. Available: https://bengrosser.com/projects/minus/ (accessed 10 May 2022).

Hall S (2019) Rethinking the 'base and superstructure' metaphor [1977]. In Morley D (ed) *Essential Essays, Volume 1: Foundations of Cultural Studies*. New York: Duke University Press, 143–171.

Helmond A (2015) The platformization of the web: making web data platform ready. *Social Media + Society* 1(2): 1–11.

Helmond A, Nieborg DB and van der Vlist FN (2017) The political economy of social data: a historical analysis of platform-industry partnerships. *Proceedings of the 8th International Conference on Social Media & Society* 38(July): 1–5.

Hind S (2019) Digital navigation and the driving-machine: supervision, calculation, optimization, and recognition. *Mobilities* 14(4): 401–417.

Hotz G (2021) George Hotz announcing COMMA THREE at COMMA_CON. Available: https://www.youtube.com/watch?v=hbLiehrC2DQ (accessed 10 May 2022).

Iliadis A and Russo F (2016) Critical data studies: An introduction. *Big Data & Society* 3(2): 1–7.

Jaton F (2021) Assessing biases, relaxing moralism: On ground-truthing practices in machine learning design and application. *Big Data & Society* 8(1): 1–15.

Katzenbach C (2021) 'AI will fix this' – The technical, discursive, and political turn to AI in governing communication. *Big Data & Society* 8(2): 1–8.

Kitchin R and Dodge M (2007) Rethinking maps. *Progress in Human Geography* 31(3): 331–344.

Langlois G and Elmer G (2013) The research politics of social media platforms. *Culture Machine* 14: 1–17.

Levy S (2020) *Facebook: The inside Story*. New York: Penguin Random House.

Luitse D and Denkena W (2021) The great Transformer: examining the role of large language models in the political economy of AI. *Big Data & Society* 8(2): 1–14.

McQuillan D (2021) Post-humanism, mutual aid. In Verdegem P (ed) *AI for Everyone? Critical Perspectives*. London: University of Westminster Press, 67–83.

Malik M and Malik MM (2022) Critical technical awakenings. *Journal of Social Computing* 2(4): 365–384.

Madrigal AC (2017) Inside Waymo's secret world for training self-driving cars. Available: https://www.theatlantic.com/technology/archive/2017/08/inside-waymos-secret-testing-and-simulation-facilities/537648/ (accessed 10 May 2022).

Masís J (2014) Making AI philosophical again: on Philip E. Agre's legacy. *Continent* 4(1): 58–70.

Merriam-Webster (2021) Cynical. Available: https://www.merriam-webster.com/dictionary/cynical (accessed 10 May 2022).

Metz C and Wakabayashi D (2020) Google researcher says she was fired over paper highlighting bias in AI. Available: https://www.nytimes.com/2020/12/03/technology/google-researcher-timnit-gebru.html (accessed 10 May 2022).

Milmo D (2021) How losing a friend to misinformation drove Facebook whistleblower. Available: https://www.theguardian.com/technology/2021/oct/04/how-friend-lost-to-misinformation-drove-facebook-whistleblower-frances-haugen (accessed 10 May 2022).

Mohamed S, Png M-T and Isaac W (2020) Decolonial AI: decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology* 33: 659–684.

Morozov E (2013) *To Save Everything, Click Here: The Folly of Technological Solutionism*. New York: PublicAffairs.

Munn L (2020) Red territory: forging infrastructural power. *Territory, Politics, Governance* 0(0): 1–20.

Nieborg D and Helmond A (2018) The political economy of Facebook's platformization in the mobile ecosystem: Facebook Messenger as a platform instance. *Media, Culture & Society* 41(2): 196–218.

Norman DA (1993) Cognition in the head and in the world: An introduction to the special issue on situated action. *Cognitive Science* 17(1): 1–6.

Phan T, Goldenfein J, Mann M, et al. (2022) Economies of virtue: the circulation of 'ethics' in big tech. *Science as Culture* 31(1): 121–135.

Plantin JC, Lagoze C, Edwards PN, et al. (2018) Infrastructure studies meet platform studies in the age of Google and Facebook. *New Media & Society* 20(1): 293–310.

Ribes D, Hoffman AS, Slota SC, et al (2019) The logic of domains. *Social Studies of Science* 49(3): 281–309.

Schäfer H (2021) Building a super human driving agent. Available: https://www.youtube.com/watch?v=tZgfDVGXdmk (accessed 10 May 2022).

Schön D (1983) *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books.

Schwartz O (2018) Competing visions for AI: Turing, Licklider and generative literature. *Digital Culture & Society* 4(1): 87–105.

Simon H (1973) The structure of ill-structured problems. *Artificial Intelligence* 4(3–4): 181–201.

Snodgrass E and Soon W (2019) API practices and paradigms: exploring the protocological parameters of APIs as key facilitators of sociotechnical forms of exchange. *First Monday* 24(2).

Sprenger F (2018) The network is not the territory: on capturing mobile media. *New Media and Society* 21(1): 77–96.

Suchman L (1992) Technologies of accountability: of lizards and aeroplanes. In: Button G (ed) *Technology in Working Order: Studies of Work, Interaction, and Technology*. London: Routledge, 113–126.

Turner F (2006) *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Chicago, IL: Chicago University Press.

Turner F (2017) *Don't Be Evil: Fred Turner on Utopias, Frontiers, and Brogrammers*. Available: https://logicmag.io/justice/fred-turner-dont-be-evil/ (accessed 10 May 2022).

van der Vlist FN, Helmond A, Burkhardt M, et al. (2022) API governance: The case of Facebook's evolution. *Social Media + Society* 8(2): 1–24.

van Dijck J (2014) Datafication, dataism and dataveillance: Big Data between scientific paradigm and ideology. *Surveillance & Society* 12(2): 1–12.

van Doorn N and Badger A (2020) Platform capitalism's hidden abode: producing data assets in the gig economy. *Antipode* 52(5): 1475–1495.

Van Loo R (2020) The new gatekeepers: Private firms as public enforcers. *Virginia Law Review* 106: 467–552.

Veale M and Binns R (2017) Fairer machine learning in the real world: mitigating discrimination without collecting sensitive data. *Big Data & Society* 4(2): 1–17.

Verdegem P (2021) Introduction: why we need critical perspectives on AI. In Verdegem P (ed) *AI for Everyone? Critical Perspectives*. London: University of Westminster Press, 1–18.

Wall Street Journal (2021) *The Facebook Files: A Wall Street Journal Investigation*. Available: https://www.wsj.com/articles/the-facebook-files-11631713039 (accessed 10 May 2022).

Waldman K (2020) Has self-awareness gone too far in fiction? Available: https://www.newyorker.com/books/under-review/has-self-awareness-gone-too-far-in-fiction (accessed 10 May 2022).

Waymo (2021) We're building the World's Most Experienced Driver. Available: https://waymo.com/ (accessed 10 May 2022).

Wu AX and Taneja H (2020) Platform enclosure of human behavior and its measurement: Using behavioral trace data against platform episteme. *New Media & Society* 23(9): 2650–2667.

YoHa (2021) YoHa. Available: http://yoha.co.uk (accessed 10 May 2022).

Author Biographies

Sam Hind is a media scholar. His research interests include navigational technologies, sensing and the sensor society, AI and algorithmic decision-making, automotive cultures, and mobile play. He is currently researching autonomous driving, machine vision 'challenges' and 'sensor work', the platformization of automobility, and the datafication of driving.

Tatjana Seitz is a researcher at the CRC 1187 "Media of Cooperation" and PhD candidate at the University of Siegen. Her research focuses on API studies at the intersection of economic, aesthetic, and data driven concepts within the context of networked social interfaces. She is equally interested in technically informed critical concepts and methodologies for the study of computational cultures.