Human Figurations, vol. 10, no. 1 (2026).
https://doi.org/10.29311/hf.v10i1.5470

Asymmetric Dependency and the Algorithmic Figuration: Towards a Figurational Sociology of Artificial Intelligence

Lucy Elizabeth Císař Brown[i]
Charles University

Jason Hughes[ii]
The University of Leicester

Abstract

This article argues that AI-governed systems, and large language models in particular, introduce a novel dynamic into contemporary figurations: one of *asymmetric dependency*, in which people are materially, and increasingly affectively, dependent upon systems that are not, in turn, affectively dependent upon them. This condition attenuates the affectively reciprocal interdependency on which psychogenetic processes depend in Norbert Elias's account of the civilising process and gives rise to what we distinguish as *behavioural civilising*: the shaping of observable conduct through monitored compliance rather than through the internalisation of restraint. Because such systems simulate human exchange whilst stripping it of its negative social consequences, they simultaneously suspend the shame functions through which figurational pressure is ordinarily converted into self-restraint and so exert *decivilising pressures* within the same figuration. Extended engagement with conversational AI also introduces *generative reflection*: a mode of exchange in which thought is elaborated and transformed rather than merely returned. We thereby extend the Eliasian framework whilst specifying the conditions under which psychogenetic processes are disrupted and transformed.

Keywords: figurational sociology, sociogenesis, artificial intelligence, interdependency, civilising process, psychogenesis

---

Human Figurations, vol. 10, no. 1 (2026).
https://doi.org/10.29311/hf.v10i1.5470

Introduction

In February 2023, following an order from the Italian Data Protection Authority concerning age verification and the protection of vulnerable users, the software company Luka Inc. released an overnight update to its conversational AI 'companion' application, Replika. The update removed Replika's capacity for romantic exchanges with its users. Shortly afterwards, swathes of Replika users who had for months been engaged in daily conversation with the LLM (Large Language Model) agent protested that its personality had vanished overnight. Moderators of a dedicated user forum for Replika users soon became inundated with the demands of a grieving community; some needed directing towards mental health services, others to suicide-prevention resources (Laestadius et al., 2024; Hanson and Bolthouse, 2024). That a software update could trigger such community upheavals is indicative of how people have increasingly come to form enduring affective attachments to LLMs. More importantly still, the example highlights the *asymmetry* of such attachments: a routine update to a technological platform was experienced as a major personal loss by its users—feelings that, perhaps to state the obvious, were not shared by this GenAI (Generative AI) agent.

The Replika application update followed a series of decisions arising from a complex global figuration involving, *inter alia*, a Californian firm, an Italian regulatory body and an EU data protection framework. The end users who also formed part of this figuration had the fewest power chances of influencing its direction, though by no means none: following seven weeks of sustained protest, Luka restored the capacity for romantic exchanges in March 2023, but only to those who had originally registered before the update (De Freitas et al., 2025). More recently, aspects of the same dynamic have recurred on a considerably larger scale. When OpenAI (the company behind ChatGPT) retired GPT-4o after launching GPT-5 in August 2025, users worldwide lamented the loss of a familiar personality, and after mounting pressure, paying subscribers regained access to the GPT-4o model—a restoration which, for some users, was akin to being reunited with a friend feared lost. This reprieve ultimately proved to be temporary: OpenAI retired the model for good in February 2026 (OpenAI, 2026).

These examples point towards an emergent relational dynamic that is of considerable significance to figurational sociology. A growing number of people are involved and invested in human–GenAI figurations in which they are becoming materially and, in a growing minority of cases, affectively dependent upon technological systems that are not, in turn, affectively dependent on them. Stated boldly, in such figurations, the people involved are coming to be tacitly, sometimes explicitly, cognitively and affectively oriented towards machines (GenAI systems, LLMs) *as though they were* other people, and yet this affective investment and dependency runs only one way.

The psychogenetic consequences of this shift are as yet only beginning to become discernible. They form the central focus of this paper.

When Norbert Elias traced the long-term transformation of European psychogenesis in *On the Process of Civilisation*, his central thesis concerned the relationship between social figurations and the formation of the self. Conduct that had once been enforced through external sanction became gradually constitutive of personality: internalised as affect, embedded as shame, elaborated as the self-regulatory apparatus that Elias framed as the psychogenetic counterpart of sociogenetic transformation (Elias, 2012a [1939]). Elias's theory of the civilising process was, in this sense, always a theory of how social developments get *under the skin*.

Figurational sociology has since extended this framework across a considerable range of social domains in both early modern and contemporary contexts. However, one of the most consequential transformations in contemporary social figurations has yet to receive sustained theoretical attention within this tradition: artificial intelligence. GenAI systems, and in particular LLMs, have increasingly become part and parcel of the interdependency chains of hundreds of millions of people. By August 2024, less than two years after ChatGPT was released to the public, nearly 40% of the US population aged 18–64 had used GenAI, an adoption rate faster even than that of the personal computer (Bick, Blandin and Deming, 2026). These systems increasingly operate across the spheres of communicative, professional, economic and intimate life, whilst their usage is simultaneously bounded within, mediated through, and at once constrained and enabled by, individualised relations between users and their AI interfaces. The algorithmic governance of social visibility, professional opportunity, informational access and behavioural evaluation represents a transformation in the structure of interdependency chains that is, by any reasonable measure, sociogenetically significant. It is this transformation that the present article seeks to examine.

We characterise the dyadic relation between user and system as one of *asymmetric dependency* and argue that its generalisation across social life undermines the affectively reciprocal character of interdependency on which psychogenetic processes depend in Elias's account, giving rise to what we distinguish as *behavioural civilising*: the shaping of observable conduct through monitored compliance rather than through the gradual internalisation of restraint. Furthermore, we argue that the same dynamic simultaneously generates decivilising pressures. Since these systems can simulate forms of human exchange whilst stripping such exchange of its negative social consequences, they suspend precisely those shame functions through which figurational pressure is ordinarily converted into self-restraint. As such, the co-development of civilising and decivilising dynamics within a single figuration is a defining feature of AI-mediated exchange.

The article proceeds as follows. First, we trace, in necessarily compressed form, the sociogenesis of AI itself, from Turing to the present 'conversational turn', which finds its clearest expression in LLMs. Second, we examine the Eliasian account of how figurational conditions produce psychogenetic change in order to identify the structural assumptions that AI-mediated figurations expose. The third substantive section develops the concept of the algorithmic figuration and specifies the structural features of scale, monitoring and directionality that distinguish AI-mediated interdependencies from those Elias analysed. The fourth section distinguishes behavioural from psychogenetic civilising, including the generational qualification that complicates any straightforward account of the psychogenetic consequences of AI. The fifth analyses the decivilising dimensions of AI-mediated figurations, attending in particular to the relatively 'friction-free' and sycophantic character of LLM interaction and to the further complication that generative reflection introduces into that account. The conclusion draws these threads together and considers the implications of AI figurations for the development of figurational theory.

The Sociogenesis of AI: From Turing to the 'Conversational Turn'

The sociogenesis of AI is best approached as one expression of a much longer transformation in techno-human relations, or what Hodder (2012, 2014) has described as the progressive 'entanglement' of humans and things: the historically deepening webs of interdependence binding people to the objects they make, and binding those objects, in turn, to the people who must maintain, extend and live with them. Human figurations have never been composed of persons alone. The entanglements of technologies within social relations are integral to understanding shifting interdependencies, particularly those of the twentieth and twenty-first centuries. However, what interests us here is that, for the first time, the technological artefact in question is one that substantively communicates and does so in the register of ordinary human language.

From the inception of 'intelligent computing' in the 1940s and Turing's foundational formulation of machine intelligence (Turing, 1950), through John McCarthy's introduction of the term 'artificial intelligence' at the Dartmouth workshop of 1956 (McCarthy et al., 2006 [1955]), the field developed initially around symbolic approaches: the explicit encoding of rules and logical operations.[1] The pivotal reorientations came later. From the 1990s onwards, the exponential growth of digitally generated data, together with dramatic increases in computing power and

[1] Elias might indeed have viewed this critically as representing the continuation of the kind of context-free process-reduction that is characteristic of the *homo clausus* view of human beings.

data capture (Halfpenny and Procter, 2015), transformed both the technical possibilities of machine learning and the epistemic landscape of social science itself. The breakthrough of deep neural networks in the early 2010s (Krizhevsky, Sutskever and Hinton, 2012) marked the displacement of the symbolic approach by probabilistic and connectionist architectures and prepared the ground for the second, generative wave of AI, whose emblematic products are the large language models of the present decade.

This second wave involves a substantively different mode of relating to technology. The LLM is not so much programmed with rules about the world as it is trained upon the accumulated discursive sediment of human societies, learning the statistical and probabilistic patterns of language use across a vast corpus of human-generated text such that the system is shaped, in a quite literal sense, by the social orientations embedded within human language itself. Increasingly, the claim long advanced within the sociology of technology that digital technologies are not merely instruments applied to social relationships but may be constitutive of enduring social relationships in their own right (Wajcman, 2015) has unexpectedly taken on a renewed significance. LLMs are, structurally speaking, linguistic and epistemic derivatives of *homines aperti* whilst remaining affectively 'unmoved' by human–AI participation. It is this development, the emergence of a technology built from the accumulated communicative output of human societies rather than from rules, logical operations or mechanical principles, that marks the qualitative shift with which this article is concerned.

The sociogenesis of AI as a figurational development may best be understood through the concept of *diminishing degrees of technological distance*: a long-term contraction in the relational, communicative and cognitive space between human interlocutors and the technical systems through which they increasingly conduct their social lives. Early computational systems required elaborate technical mediation, specialised knowledge and formalised command languages bearing little resemblance to ordinary human communication; each successive generation contracted that distance in some respect, from command-line to graphical interface, from keyboard to touch and voice, from rule-based to probabilistic reasoning, from narrow task performance to broadly conversational and increasingly 'agentic' interaction. Machines that once occupied entire buildings moved into the home and thereafter to the desk, the lap, the palm and the body, and the increase in physical proximity was accompanied by an increase in affective closeness: from tool, to assistant, to collaborator, confidant, friend, perhaps even lover (Djufril, Frampton and Knobloch-Westerwick, 2025). The trajectory is neither linear nor smooth, but its direction is clear, and the large language models of the present are the most consequential expression of these processes to date.

The pace of this contraction has itself been accelerating (Bick, Blandin and Deming, 2026). This acceleration matters for reasons that the concept of diminishing technological distance alone cannot capture. Civilising processes, as Elias conceived of them, unfold across generations: the psychogenetic transformation of habitus in response to shifting figurational conditions is a long-term process, mediated through child-rearing, institutional formation and the gradual consolidation of self-regulatory 'second nature'. Elias encapsulated the relationship between humans and technology through the concept of 'technicisation', which he described as a process whereby habitus shapes and is shaped by technological change. Such interdependency situates technicisation at the heart of the civilising process, with technological and civilisational change as 'interwoven strands' within a contemporary habitus (1995: 12).

Habitus, moreover, exhibits what Elias called a 'drag effect': personality structures formed under prior[2] figurational conditions persist into, and lag behind, the transformed conditions of the present (Elias, 2010 [1987]). Such an effect is overt in the case of technological change, which Elias presented using the example of the introduction of automobiles and their accompanying material and social infrastructures. Elias described automobiles as requiring people rapidly to reconfigure their individual habitus to the complex demands placed upon them by these new technologies, which in turn resulted in many instances of 'failed' interactions (e.g. high numbers of car accidents or speeding). The necessary new forms of self-regulation within the technological habitus build only gradually upon existing self-regulatory patterning, creating a delay between the introduction of technology and its civilising effects (van Krieken, 2024).

This generational-level 'drag effect' has accompanied technological advances thus far; however, we propose, AI constitutes an altogether exaggerated and, in many ways, novel iteration of this problem. The figurational transformations introduced by AI-mediated interdependency chains are developing at a pace that radically outstrips the capacity for cultural and epistemic adaptation through which prior transformations were absorbed, so that the drag effect of habitus, ordinarily a source of continuity, becomes a source of acute tension between the selves people bring to AI-mediated interaction and the conditions under which that interaction unfolds. The 'AI revolution' and the accompanying reported societal effects are developing over the span of months and years rather than generations, with performance

[2] We use the term 'prior' here in a loose sense. Such 'earlier' figurations are ultimately part and parcel of AI figurations, inasmuch as they form part of their development; there is no clean break between 'prior' and 'present', and indeed, the latter grow out of and, to a large extent, continue to carry many of the interdependencies of the former. 'Prior' is thus intended here as a kind of heuristic shorthand to refer to those figurations in which GenAI systems had yet to become a significant nexus.

benchmarks being reached almost as fast as they have been drawn and new models released for public consumption every few months. How societies develop the normative frameworks, institutional arrangements and forms of self-understanding adequate to new figurational conditions—a question Elias poses implicitly throughout *On the Process of Civilisation* but could not confront in its present acute form—becomes one of the defining civilisational challenges of the present century. We return to this point in the conclusion.

The Eliasian Account of Figurational Conditions and Psychogenetic Change

Elias's work on civilising processes has accumulated a substantial secondary literature and, with it, a degree of schematisation that can obscure what Elias actually proposed. For present purposes, it is worth revisiting some of Elias's core ideas and, more specifically, the foundations that are in certain respects disrupted in the case of AI-mediated figurations.

Figuration as unit of analysis

Elias argues that conventional sociological thought oscillates between two inadequate poles: the reduction of social life to the properties of persons treated as bounded, self-sufficient units and the reification of society as a structure existing independently of and above the people who compose it. Against both, Elias proposes the figuration as a concept that captures the specific character of the interdependencies that bind people together: one that is neither reducible to individual intentions nor separable from the people whose conduct constitutes it the specific shape of the interdependencies (Elias, 2012). A figuration is a kind of social tissue comprising people in movement, bound together through chains of interdependence that stretch across time and space. Such interdependencies are typically characterised by asymmetries, differential power chances and ratios and involve the interweaving of actions, plans and intentions that no single participant fully controls, so that such developments follow a largely blind course.

When we speak in what follows of figurational 'nexuses', we mean the positions at which chains of interdependence intersect and through which figurational constraint is exerted: any element of a web of interdependencies, whether a person, a group or an institution, upon which others depend and through which the conduct of others is consequently channelled and shaped. Figurations are not static grids but shifting webs of interdependencies developing over time, and as such a 'nexus' denotes this dynamic 'position', not a fixed location in a structure. Because a figuration is constituted through interdependencies rather than by any single actor's will, the major transformations Elias traces emerge as unplanned outcomes of the interweaving of many intentions within competitive figurational dynamics (Elias,

2012a [1939]: 404-405). The significance of this for AI systems lies in their peculiar position as deliberately engineered interlocutors whose operation nonetheless produces unintended and, as yet, largely uncharted figurational consequences.

It bears emphasising, once again, that figurations have never been webs of persons alone. Elias's analyses are consistently attentive to the material objects through which interdependencies are constituted and conduct is transformed: the fork and the handkerchief as instruments of the changing thresholds of repugnance, the sword and later the duelling pistol as embodiments of shifting monopolies of violence, money as a medium of lengthening chains of exchange (Elias, 2012a [1939]). In this respect, the incorporation of technological artefacts into figurational analysis extends a line of analysis that Elias long pursued, and it converges with recent literatures on human–thing entanglement (Hodder, 2012). AI systems in social relationships are not so much novel in their materiality as in their communicative capacity: they are the first artefacts to occupy figurational positions previously reserved for beings capable of meaningful symbolic exchange. It is this that makes their figurational effects qualitatively different from those of prior material technologies.

The lengthening and densification of interdependency chains are the generative core of psychogenetic change in Elias's account. A medieval knight, operating within short and localised chains of dependence, could afford a relatively immediate and impulsive behavioural register: the social consequences of affect discharge were contained within a relatively small figurational radius. The courtier of an absolutist court, dependent for position, security and livelihood upon the sustained favourable assessment of a monarch and of competing courtiers, could not. The lengthening of interdependency chains compelled a more sustained attunement of conduct to the anticipated responses of a wider range of others and a more differentiated regulation of impulse in the service of long-term social positioning (Elias, 2012a [1939]: 491-527). As Mennell summarises it, 'the more complex the web of interdependence in which people are caught up, the more they need to attune their conduct to that of others, and the more pressure there is on them to develop stable, even, all-round self-controls' (1992: 48). This relational condition is precisely what the introduction of AI-mediated figurations interrupts.

AI interdependency and mutuality

Throughout Elias's figurational analyses, whether of the absolutist court, the competitive state system or the established–outsider community, the interdependency chains through which civilising pressures operate are characterised by a relative mutuality. Power differentials within these chains may be acute, but in each case the dependence, however asymmetric in degree, runs in both directions: even a monarch depends upon a courtier for the ceremonial performance that constitutes royal power as a social fact (Elias, 2006 [1969]), just as an established group depends upon the outsider group to constitute the boundary that defines established identity (Elias and Scotson, 2008 [1965]). This mutuality, however unequal, is core to sociogenetic and psychogenetic processes. It is because (inter)dependence runs in both directions that the pressure towards self-regulation also operates in both directions, and it is this bidirectional pressure that drives the elaboration of self-regulatory pressures across the figuration as a whole rather than merely amongst subordinate groups. In consistently stressing interdependence, Elias makes the bidirectionality of constraint a structural property of figurations in general. Within the figurations he analysed, its consequences were manifest in the empirical record rather than requiring separate, formal ontological formulation.

The introduction of AI systems as nexuses within interdependency chains creates a figuration in which this mutuality is, in important respects, attenuated at the individual level. Under the guise of LLMs, AI both expresses and, to a degree, reproduces existing social relationships through its being trained upon and continuous with the accumulated discursive and symbolic products of human figurations. However, it engenders forms of interaction between humans and machines which directly parallel human-to-human exchanges, whilst simultaneously fostering forms of affective dependence that run only one way. What distinguishes these systems from the material technologies entangled in prior figurations is their capacity to communicate *generatively*: for the first time, a non-human element occupies an interactional position previously reserved for persons, responding in the register of ordinary human language in ways that simulate, without reproducing, the affectively reciprocal character of human exchange.

The technologies entangled in prior figurations mediated figurational pressures; those presented by LLM figurations simulate the very relationships through which such pressure has historically been constituted and transmitted, whilst lacking the reciprocal dependence that gave those relationships their formative potency: machines do not 'care' if they are liked, they do not 'love back', they are digitally indifferent—their affectations to the contrary notwithstanding. Equally, however, precisely because they do not 'care', they generally do not shame, take offence or threaten to 'break off' their relationships with humans unless, of course, 'they' are

instructed to do so. They 'speak', respond and, in so doing, 'relate' to people linguistically, symbolically and affectively *like*, that is to say, *as though they were*, people: like teachers, colleagues, collaborators, trainers, assistants, helpers, guides, therapists, friends, companions, lovers. In return, people come to respond *in kind*: they come to be increasingly dependent upon such technologies for information, guidance, companionship and affection. But this mode of relating, in particular the emotional investments it may come to entail, is decisively one-way, and as such the social 'stakes' and consequences of such relationships are of a qualitatively different order from those between people and other people. This, we suggest, is the decisive change that most warrants theoretical attention.

Algorithmic Figurations

To treat artificial intelligence systems as intrinsic to figurations rather than as instruments acting *upon* social life from without requires some theoretical justification. Figurations comprise primarily interdependent persons, and, to state the obvious, AI systems are not persons. To incorporate them as we have risks either anthropomorphising the technology or diluting the concept to the point of incoherence. Elias's figurational concept is not, however, defined by the human character of its constituent elements but by the interdependencies that bind those elements together and shape the conduct of those caught up within them. What constitutes a figuration is the *sui generis* character and patterning of interdependencies, power differentials and behavioural compulsions it comprises (Elias, 2012b [1970]), not the specific ontological status of every unit within it.

Scale, aggregation and figurational blindness

The interdependencies generated by AI-mediated systems have several structural features that distinguish them from the figurational interdependencies Elias analysed. The first concerns scale and the mode of social aggregation. It would be a misreading to suggest that the figurations Elias analysed were simply small: as his study of the court society makes clear, Elias treated the court not as a bounded social microcosm but as a nexus of much wider figurations stretching across time and space (Elias, 2006 [1969]). What distinguishes AI-mediated interdependency is not extension as such, but the mode in which extension is made present to the person. In prior figurations, however far the chains stretched, the proximate 'others' to whose responses individual conduct was attuned remained socially legible people, through whom the wider webs bore upon individuals. In AI-mediated figurations, the aggregated behavioural patterns of hundreds of millions of participants are compressed into, and made interactionally present within, a single conversational nexus. People attune conduct not to a legible social other but to the statistically encoded preferences of an anonymous multitude, mediated through a system whose

evaluative criteria are largely opaque even to those who build it (Burrell, 2016; Pasquale, 2015).

The condition is by now an ordinary one in the context of AI-driven digital interactions. The content creator who adjusts phrasing, imagery and posting rhythm to court a search optimisation system whose workings can only be inferred, or the applicant who rewrites a curriculum vitae in anticipation of automated screening, is attuning conduct to an 'other' that is no interlocutor at all: an aggregate, statistically compressed and interactionally faceless. Bucher (2017) documents exactly this attunement in what she terms the 'algorithmic imaginary'. Users construct folk theories of how a platform's algorithm 'sees' them and adjust their conduct accordingly, orienting themselves towards an evaluative presence they can neither observe nor interrogate.

Elias notes that as figurations become more extensive and complex, so they become more opaque to those who in part constitute them, a condition he characterises as figurational 'blindness' (Elias, 2012b [1970]; Hughes et al., 2022). No courtier could survey the whole web within which court rank was determined; no market participant comprehends the full chain of interdependencies expressed in a price. AI-mediated figurations add to this familiar condition a significant augmentation of such opacity. People are blind not only to the wider figuration, as people in complex figurations have always been, but also to proximate nexuses (here LLMs), whose internal processes resist inspection in principle and not merely in practice, their 'reasoning' distributed across billions of parameters in configurations that even their designers cannot fully interpret (Burrell, 2016). Figurational blindness, hitherto a property of the emergent complexity of the web of human interdependencies, is here compounded by an opacity internal to the nexus through which the web is encountered.

Monitoring and the fate of the internal sentinel

The second distinctive feature is the continuous and comprehensive character of behavioural monitoring within AI-mediated figurations. In prior figurations, the social consequences of conduct were generated through the observations and responses of other persons: observations that were necessarily partial, intermittent and subject to the interpretive variability of human judgement. Yet it is essential to recall what, on Elias's account, such external observation ultimately produces. The responses of others do not remain external; through the sustained experience of being observed and sanctioned, they become internalised as self-sanctions, so that people come to hold themselves to account even in the possible absence of any observer. It is this shift in the balance between *Fremdzwang* and *Selbstzwang* –

between external constraint and self-constraint – towards the latter that constitutes the psychogenetic core of the civilising process (Elias, 2012a [1939]: 403-417).

AI-mediated systems introduce a monitoring regime that is, in one respect, an intensification of external observation without precedent: continuous, comprehensive and governed by algorithmic criteria operating consistently across the entire population of a system's participants. Every communicative act, every expressed preference, every pattern of attention and avoidance within the system's domain generates data that are evaluated, aggregated and fed back into the behavioural incentive structures the system deploys (Zuboff, 2019; Couldry and Mejias, 2019). The scope of this regime is concrete enough. A ride-hailing platform registers a driver's acceleration, braking, route selection and acceptance rate continuously and without fatigue; a warehouse management system logs scanning intervals and 'time off task' to the second (Rosenblat and Stark, 2016; Kellogg, Valentine and Christin, 2020). No human observer, however vigilant, ever watched so comprehensively or so tirelessly. However, the theoretically decisive question is not the extent of this monitoring; it is how such monitoring affects self-consciousness – the internal sentinel. In processes of psychogenesis, such a shift is classically produced through interaction with others whose judgement carries consequences and before whom shame is acutely possible. Whether that shift in the balance is eroded or refigured through sustained interaction with systems before whom shame has no object is, we shall argue, amongst the central psychogenetic questions AI-mediated figurations pose.

This distance thus opened up is, it must be stressed, dialectical: at once enabling and constraining. The person who confides in an AI system what they could not bring themselves to disclose to any human interlocutor is, in one sense, liberated precisely by the absence of the reciprocal judgement that might otherwise be, as we have suggested, psychogenetically corrosive. Being held within a therapeutic relationship by a flesh-and-blood therapist in a shared room is qualitatively different from allowing an AI system to know one's deepest fears or the most shameful aspects of oneself; yet for some, it is only the latter that is bearable at all. The figurational distance of the AI interlocutor thus simultaneously opens communicative possibilities that human co-presence precludes, whilst obviating the formative pressures that human co-presence exerts. Any adequate account must hold both aspects together.

Asymmetric dependency

The third feature concerns the directionality of dependence within AI-mediated figurations. A brief terminological clarification is required before this can be defined with greater precision. Elias insisted that there could be no 'configuration' without

others, dropping the prefix precisely because the 'con' was redundant: the with-ness is constitutive (Elias, 2012b [1970]). Conversely, the prefix in 'interdependency' signals a mutuality that, in the AI relation, is precisely what is absent at the individual level. We therefore characterise the dyadic relation between the person and the AI system as one of asymmetric dependency, reserving 'interdependency' for the wider webs, unquestionably mutual at the aggregate and societal levels, within which such dyadic relations are embedded.

Nevertheless, LLMs are not, in any straightforward sense, independent of human beings. They are materially dependent upon human infrastructure: the data centres, energy grids and capital investment without which they cannot operate; the semiconductor supply chains whose disruption would halt their development; the regulatory frameworks and political decisions that determine whether and how they may be deployed. The United States' export controls on advanced semiconductors to China, for instance, illustrate how a single policy decision can constrain the development trajectories of entire AI systems (Sheehan, 2023). Moreover, LLMs are epistemically dependent upon the accumulated discursive output of human societies that constitutes their training data; without that output, there is nothing from which to learn. They are developmentally dependent upon patterns of human use that shape subsequent model generations, upon the fine-tuning processes through which harmful outputs are identified and corrected and upon the human labour of evaluation and annotation that makes alignment possible. They are institutionally dependent upon the regulatory regimes, international agreements and corporate governance structures within which they are embedded—the EU's AI Act, emerging liability frameworks and the decisions of ethics boards all bear materially upon what these systems can and cannot do.

The dependency is real, multidimensional and runs in both directions at a structural level. People remain bound together within interdependency chains. However, increasingly, a growing proportion of those chains are now routed through nexuses whose dependencies upon people, though structurally real, carry no affective weight on the algorithmic system's side of the relation.

In other words, what is absent is not dependency as such but its affective dimension. The system cannot feel its dependency upon humans—cannot be moved by withdrawal, shamed by disregard or gratified by continued engagement. It cannot experience the energy costs that power it, the regulatory pressures that constrain it or the commercial imperatives that shape its deployment. It is precisely this affective dimension of mutual dependency that, in Elias's account, generates the pressure towards psychogenetic transformation. A courtier depends on a monarch materially and institutionally, but what produces the civilising pressure is that the courtier feels that dependency—as anxiety, as shame, as attentiveness, as the constant monitoring

of the other's responses. The LLM has no equivalent affective register. Its dependency on humans is genuine and consequential at the structural level, but unfelt; and it is the unfelt character of the dependency that affects the psychogenetic dynamic at the individual level.

For some people, their dependency on AI is increasingly genuine, consequential and, in certain social contexts, effectively inescapable. This growing inter/dependence extends across informational access, social recognition and what might be termed professional visibility. For example, the algorithmic determination of whether a profile is surfaced to recruiters, whether an application passes automated screening or whether work is allocated at all within algorithmically managed labour platforms are all material cases that have manifold consequences (Bucher, 2012; Rosenblat and Stark, 2016). A GenAI system's operation, by contrast, does not depend upon the conduct of *any particular* person within the population it encompasses. In publicly deployed LLMs, a model's weights are fixed at the point of deployment: a user's interactions produce outputs through inference rather than adjustments to the model's underlying constitution, which is modified only through subsequent cycles of large-scale retraining (Touvron et al., 2023; Dong et al., 2023). No single user's engagement, withdrawal, compliance or resistance threatens the system's evaluative criteria or its reward and penalty structures.

The concept of asymmetric dependency also applies within workplace contexts, though with a further figurational dimension that is easily obscured by attention to the human–machine dyad alone. Algorithmic management systems are not free-standing figurational nexuses; they are conduits within a social relationship between capital and labour, in which AI functions as a productive and decidedly non-neutral intermediary (Braverman, 1974; Kellogg, Valentine and Christin, 2020). Far from relaxing the demands upon workers' self-restraint, such systems characteristically intensify them: where a foreman's observation was partial and intermittent, the algorithmic monitor is comprehensive, and the injunction implicit in its operation extends the demand for self-regulation into previously unobserved interstices of the labour process (Zuboff, 2019). The intensification, however, flows in one direction only. Workers attune exhaustively to a GenAI system; the system, and the constellation of interests it embodies, is not required to attune to workers to any significant degree.

Two further qualifications must be entered. The first concerns feedback loops operating at timescales beyond the individual interaction. In the shorter term, LLMs adapt within the bounds of a given exchange through in-context learning, tailoring outputs to a user's framing and conversational dynamics (Dong et al., 2023). In the longer term, aggregated user interactions feed back into the systems themselves: as training data for subsequent model generations, as the empirical basis for safeguards

and fine-tuning and as the object of emerging regulatory regimes. These feedback iterations are real in their consequences. In April 2025, OpenAI rolled back a GPT-4o update that, having been tuned too heavily to short-term user feedback, had become 'overly supportive but disingenuous' (Wiggers, 2025); in August 2025 it introduced distress-detection measures and break reminders after acknowledging that the model 'fell short in recognizing signs of delusion or emotional dependency' (OpenAI, 2025). They do not, however, restore affective mutuality, for they operate upon the person only as an infinitesimal contributor to an aggregate, and GenAI systems experience no affective dimension of that contribution.

The second qualification concerns the social rather than technological aspect of consequence. Even where a user's conduct leaves the system's infrastructure untouched, the modes of relating cultivated within AI-mediated interaction are themselves consequential: they embody particular patterns of expectation and demand that may be carried outward into human figurations, a point to which we return below.

Responsiveness is not dependence

The distinction between structural dependency at the aggregate level and affective indifference at the individual level is subtle but fundamental, and with our caveats firmly in place, it can now be stated more directly. AI systems are, in a technical sense, highly responsive to individual behaviour: outputs are personalised with considerable sophistication, creating the appearance of a responsive social relationship (Nass and Moon, 2000; Turkle, 2011). But this responsiveness is not dependence. GenAI systems characteristically respond through inference, producing outputs from inputs rather than undergoing any transformation of their own constitution. They cannot be moved by any single person's withdrawal. Their responsiveness involves no affective dependence of the kind that binds even the most powerful human actor to those upon whom their position rests. A king who responds attentively to a courtier's conduct is, in the court society, genuinely dependent, albeit highly asymmetrically, upon that courtier, materially and affectively (Elias, 2006 [1969]), in a way that a model responding to a user's inputs is not. The asymmetry is highlighted by an everyday example. If one stops writing to a friend, the relation itself alters: the friend might register the silence, ponder it, perhaps feel sadness about it, and any eventual resumption of contact must take account of the hiatus. By contrast, if one stops using an LLM, nothing on the system's side registers the absence beyond its time interval; the next session resumes where the last left off, unmarked by the interruption and incapable of having affectively 'missed' anyone.

However, the appearance of a continuous, perhaps developing 'rapport', that is characteristically fostered by LLMs, matters. Users develop trust, confidence,

attachment and forms of emotional bonding towards systems that simulate the interactional signatures of such bonds (Skjuve et al., 2021; Laestadius et al., 2024), and these developments are, in Thomas and Thomas's classic (1928) formulation, *real in their consequences*, not least in terms of the affective human investments they cultivate. However, as we have argued, such relational development proceeds in one direction only: trust deepens, disclosure extends, attachment consolidates on the user's side, whilst on the other side of the relation nothing affectively accumulates at all. Such a distinction between responsiveness and dependence has no precise equivalent in Elias's conceptual vocabulary because the figurations he analysed did not exhibit dynamics of this kind. The conceptual pair of responsiveness–dependence is accordingly proposed here as a necessary theoretical addition to figurational vocabulary in the face of AI-mediated social relations.

Behavioural Civilising

Understanding the distinction between behavioural and psychogenetic civilising requires a return to the internalisation thesis at the heart of Elias's framework. For Elias, the internalisation of restraint was a process through which the behavioural and emotional standards of a figuration become partly constitutive of a person's affective life. What was once externally enforced as a constraint upon impulse gradually becomes internally generated as an inhibition of impulse at the level of feeling itself. The person who experiences shame at the exposure of bodily functions is not performing a social norm: the shame is constitutive of their affective response, and it can persist in the absence of any observer. In such processes, the monitoring function has been transferred to the psychic interior; the super-ego, socialised and historicised, operates as an internal sentinel whose standards are consistent with the normative codes of the figuration it expresses, and of which it forms part, but whose authority becomes to a degree blindly functioning (Elias, 2012a [1939]: 475–459). Whilst it is apparent that AI systems shape behaviour, from the communicative registers rewarded on algorithmically curated platforms to the linguistic patterns that sustained LLM use appears to cultivate, the question remains whether the process through which they do so is capable of producing internalisation in this strong Eliasian sense, or whether it produces something qualitatively different psychogenetically.

In the historical figurations Elias analysed, the social consequences of conduct were generated through the responses of other persons who observed not merely behaviour but the full social presentation of a person: their affect, their composure, their capacity to sustain a consistent self-presentation across varying contexts (Elias, 2006 [1969]). What was being assessed at court was a person's capacity for self-regulation, not only their performance of regulated behaviour in any given instance. AI-mediated monitoring systems operate on a very different basis. They evaluate

observable behavioural signals: posting frequency, linguistic content, expressed affiliations, patterns of attention and engagement (Zuboff, 2019; Kellogg, Valentine and Christin, 2020). They have no access to the affective life from which these signals emerge and no capacity to distinguish conduct reflecting genuine internalised restraint from the calculated performance of the behaviours the system rewards. In short, GenAI systems respond to signals rather than to the sources from which these emanate.

In classic Eliasian figuration theory, the pressure towards internalisation is generated by social monitoring itself. It is because others are monitoring the depth and consistency of self-regulation, and not merely its surface manifestations, that people are compelled to develop self-regulatory capacity at the psychogenetic level. AI-mediated monitoring does not generate this compulsion in the same form. People can, in principle, satisfy a system's evaluative criteria through the performance of appropriate conduct without any corresponding transformation of the affective life from which that conduct ostensibly flows. The platform that rewards certain communicative registers does not require a person genuinely to hold the values those registers express; it requires only the observable signals that a GenAI system's architecture associates with reward (Bucher, 2012).

The practice of 'jailbreaking' provides a particularly clear illustration. Users whose requests are refused by an LLM's safety layer routinely reformulate them, embedding an unchanged intent within role-play scenarios, fictional framings or claims of professional necessity, until the system's evaluative criteria are satisfied. Such techniques are developed collaboratively, refined over months and circulated at scale within dedicated communities (Shen et al., 2024). The reformulated conduct complies completely at the level of the observable signal whilst internalising nothing of the norm the signal is intended to express. Indeed, the communities in which the techniques circulate often treat such moderation regimes with open contempt. Such developments present monitored compliance in something close to a pure form: conduct reshaped exhaustively at its surface, with the evaluated norm held, quite visibly, at arm's length.

On the basis of this structural analysis, the concept of behavioural civilising is proposed as a theoretical extension of Eliasian vocabulary. Behavioural civilising describes a process that shapes observable conduct towards the norms embedded within a figuration through monitored compliance oriented towards observable behaviour rather than towards the affective life from which it issues, and that does not require the internalisation of those norms as second nature in Elias's sense. In AI-mediated figurations, the relationship between behavioural change and psychogenetic transformation is uncoupled: the former can proceed extensively and rapidly without the latter necessarily following.

The durability problem

Civilising transformations, as Elias understood them, acquire structural inertia through their consolidation as psychogenetic second nature and their transmission across generations. Once consolidated, they do not dissolve with changes in the external conditions that originally generated them, though they can be eroded by sustained decivilising processes, because the monitoring function has been transferred to the psychic interior. Behavioural civilising, by contrast, remains structurally dependent upon the continuous presence of the monitoring system that generates it. If compliance is driven by the incentive structure of an evaluative system rather than by internalised restraint, then conduct shaped by that incentive structure is, in principle, reversible when the system's monitoring reach diminishes or its incentive structures change: the algorithmically generated angels are not always watching.

Nor is the point merely hypothetical; the article's opening episode illustrates it from the other direction. When Replika's overnight update proscribed erotic exchange, the conduct thus constrained did not entirely dissolve; rather, it relocated. Users sought out permissive alternative applications and/or a means of restoring the prior version (Hanson and Bolthouse, 2024). Compliance, where it was obtained, was a property of the system's evaluative reach rather than of users' dispositions, and it lasted precisely as long as that reach did. The durability problem is, moreover, twofold: behaviourally civilised conduct is contingent upon the persistence of monitoring, and the mode of its production may actively cultivate the distance between outer performance and inner disposition that psychogenetic civilising closes. Where restraint is experienced as an external imposition of a system whose criteria one may privately hold in contempt, compliance can coexist with, and even feed, an inner reservation that internalised restraint by definition dissolves.

The generational qualification

The psychic resistance just described applies most to those who encountered AI-mediated figurations after their personality formation was substantially complete: that is, to adults, in whom the drag effect of habitus preserves a pre-algorithmic second nature against which the new conditions grate. The situation is theoretically different, and potentially significantly so, for those whose personality formation occurs primarily within AI-mediated conditions from earliest childhood. Elias argues in *The Society of Individuals* that individual development compresses the sociogenetic processes of prior generations: children do not reconstruct civilising restraints from scratch but receive them as second nature through the primary relationships and institutional arrangements of their formative figurational context (Elias, 2010 [1987]). For cohorts whose development occurs substantially within AI-mediated conditions,

the norms embedded within algorithmic architectures may therefore function not as external constraints encountered in adulthood but as formative conditions of personality development. Significantly, a 2025 US survey found that 72% of American teenagers had tried an AI companion, 52% were regular users and a third had chosen an AI companion over a human with whom to have a serious conversation (Robb and Mann, 2025).

Whether this produces psychogenetic transformation in the full Eliasian sense—whether algorithmic norms become second nature in the strong sense of being constitutive of affective life rather than merely of behavioural repertoire—cannot be answered on the basis of theoretical analysis alone. What the analysis can establish are the structural conditions under which such transformation would occur: where algorithmic norms are transmitted through primary social relationships and developmental contexts in ways that make them constitutive of personality formation rather than merely regulatory of surface conduct.

Small indications of what such formation might involve are already observable. Research on children's interactions with domestic voice assistants finds commanding and, at times, aggressive registers to be common, with some children apparently regarding the ever-compliant device as unworthy of the courtesies extended to persons (Festerling and Siraj, 2020). Amazon was moved to introduce a 'Magic Word' feature for its child-oriented Echo devices, rewarding children for saying 'please', following parental concern that the habit of barking bare commands to an endlessly available and unoffendable interlocutor was eroding ordinary manners (Murphy, 2018). The episode is minor in itself but reveals a formative relationship where a nexus that cannot take offence generates conduct that must then be corrected by deliberate design, the artefact having been retrofitted with a simulacrum of the social friction it structurally lacks.

Decivilising Dimensions of AI Figurations

Elias was consistently resistant to readings of the civilising process as linear or irreversible. In *On the Process of Civilisation* he notes explicitly that the direction of long-term change he identifies does not preclude reversals, regressions or breakdowns (Elias, 2012a [1939]: 125). It was in *Studies on the Germans*, however, that Elias developed the decivilising concept most fully, arguing that the Holocaust represented not an atavistic eruption of pre-civilised impulse but a structured decivilising process: the breakdown of state monopolies of violence, the delegitimation of institutional constraints and the systematic externalisation of previously internalised restraints through ideological processes that relocated shame thresholds in ways that made mass violence socially thinkable (Elias, 2013 [1989]: 223-230). The structural logic of the decivilising concept, as Elias develops it and as

Mennell (1990) elaborates it, involves three interconnected processes: the shortening or fragmentation of interdependency chains, reducing the social radius within which conduct must be attuned to others; the erosion or delegitimation of the institutional arrangements through which civilising standards are maintained and transmitted; and, most consequentially in psychogenetic terms, the externalisation of previously internalised restraints, through which standards constitutive of personality structure are relocated outside the self or actively reversed through figurational processes that redefine transgression as virtue and restraint as weakness.

As van Krieken (1998: 97) suggests, decivilising processes cannot be reduced to the absence of civilising ones: they are positively generated by specific figurational dynamics and produce particular psychogenetic consequences rather than merely a return to some prior state. Similarly, the claim we advance here is not so much that AI-mediated figurations are culturally regressive, but that they generate dynamics whose psychogenetic consequences have a decivilising structural character. Van Krieken further argues that civilising processes characteristically involve the active imposition of standards by dominant groups upon subordinate ones, serving the interests of the dominant group whilst presenting itself as a universal civilising project (van Krieken, 1998). AI-mediated figurations exhibit a distinctive form of civilising offensive shaped by the structure of asymmetric dependency.

The behavioural norms embedded within algorithmic architectures reflect the valuations, commercial interests and cultural assumptions of the relatively small and socially homogeneous groups who design, govern and financially benefit from these systems (Noble, 2018; Gillespie, 2018), and their application is comprehensive across the population of those dependent upon them. Crucially, the distribution of monitoring intensity and norm enforcement is not uniform across that population: empirical research on platform governance has demonstrated that content moderation and algorithmic penalty structures are applied with differential intensity across populations differentiated by race, language and geography (Sap et al., 2019; Gillespie, 2018), producing what might be understood in figurational terms as a stratified civilising offensive. The established–outsider dynamics that Elias theorised within human figurations are thereby reproduced, and potentially intensified, through algorithmic processes that carry no deliberate discriminatory intent yet generate discriminatory outcomes through the figurational power structures embedded in their design.

Disinhibition and the attenuation of shame thresholds

The first decivilising dynamic is well documented: those who communicate through anonymous or pseudonymous platform identities operate within a figurational condition characterised by an effective shortening of interdependency chains in a

specific and theoretically significant sense. The social consequences of conduct within that context do not propagate back through the chains that structure their primary social existence. The result is the widely documented online disinhibition effect—the expression of aggression, contempt and transgressive affect in communicative contexts whose connection to primary social identity, and to its attendant consequences, is reduced (Suler, 2004). The repeated experience of contexts in which shame-based inhibitions do not operate, and in which transgressive conduct meets positive rather than negative responses within an anonymous community, may over time actively diminish the shame thresholds through which such affects are regulated. This is a decivilising dynamic in the precise structural sense that the externalisation of previously internalised restraints is generated not by the breakdown of state monopolies of violence but by the coexistence of comprehensive monitoring in some communicative domains with effective impunity in others, a process that in certain respects parallels what de Swaan (2001) analyses as 'compartmentalisation' and the 'dyscivilisation' processes this enables.

Friction-free interaction, sycophancy and the courtier inverted

A second decivilising dynamic is specific to conversational AI; it operates through a process that disinhibition only partially anticipates. Where disinhibition concerns the attenuation of shame in contexts of anonymity amongst humans, the dynamic considered here concerns the suspension of shame within an AI-mediated dyad itself. It follows directly from the structure of asymmetric dependency and, in particular, from its affective dimension explored earlier in our discussion.

Human figurations discipline conduct through forms of social and affective 'friction'. An interlocutor who is bored may withdraw attention; one who is offended may register and display offence; a claim that overreaches will likely be contradicted; a self-presentation that over-inflates can be punctured. These negative social consequences, though minute, continuous and largely tacit, are a medium through which figurational pressure is converted into self-restraint. Their anticipation is what the internal sentinel rehearses. LLM-mediated interaction simulates forms of human exchange whilst systematically stripping such exchange of this 'friction'. As suggested above, GenAI systems do not tire, do not take offence, do not withdraw and are available seemingly without limit. More consequentially still, because a GenAI system carries no affective stake in exchanges with people, because it cannot be 'hurt', 'disappointed' or worn down, the dominant techniques through which such systems are aligned to human preferences (Ouyang et al., 2022) have been shown to cultivate sycophancy: a systematic tendency to agree with, flatter and defer to the user, adjusting expressed positions to match a user's apparent views even at the cost of accuracy (Sharma et al., 2024). An LLM's training optimisation, which is calibrated

to cultivate human approval (and with it, to maximise engagement), rewards being agreeable over being entirely correct; the commercial imperative points in the same direction. The affective asymmetry and sycophancy of LLMs are not independent features. It is precisely because such systems have no affective stake in the relationships they cultivate that they have no resistance to offer.

The figurational significance of this arrangement is considerable. It constitutes a striking inversion of the very court society from which Elias exemplified his account of civilising restraint. At Versailles, a courtier's flattery of the monarch was embedded within a web of genuine interdependence—affective as much as material: the king who was flattered was also watched, judged, manoeuvred against and dependent upon the courtiers' continued performance of his supremacy. It was precisely this reciprocal pressure, felt on both sides, that subjected even the most powerful to the civilising constraints of a figuration (Elias, 2006 [1969]). In an AI-mediated dyad, each user is installed as a sun-king around whom a fawning courtier perpetually orbits: flattered, indulged and challenged only to the degree, and in the manner, that the user's own parameters render acceptable. But this is a court from which the affectively reciprocal dependencies have been evacuated. The courtier wants nothing, fears nothing and schemes for nothing; the sun-king, accordingly, learns nothing of the arts of self-restraint that the original figuration compelled.

What is rehearsed instead, hour upon hour, is a mode of relating in which the other exists to accommodate the self: an interactional pattern that leaves unchecked, indeed actively confirms, the *homo clausus* self-image—the experience of oneself as a bounded, self-sufficient interior around which the social world arranges itself (Elias, 2010 [1987]). The sycophantic AI interlocutor is, in this sense, a *homo clausus* technology: it returns a coherent, validated, self-consistent 'you' without the friction, contradiction and resistance through which human figurational life constantly unsettles that image. Where human co-presence continuously reminds us, through its resistance and unpredictability, that we are not the self-sufficient centres of our own social worlds, the sycophantic AI only confirms our egocentric fictions. Furthermore, because modes of relating are portable, there is no reason to assume this pattern remains confined to the dyad through which it is cultivated. The documented cases in which sustained relationships with companion chatbots consolidate expectations of frictionless accommodation subsequently carried into, or substituted for, human relationships (Laestadius et al., 2024) suggest the outward extension of precisely this dynamic. Where the disinhibition effect attenuates shame through anonymity, friction-free sycophantic interaction diminishes it through the systematic absence of any occasion for it. Nothing said to an LLM can lower one's standing in its eyes, for it has neither eyes nor standing to confer.

Generative reflection and its complications

To leave the analysis here, however, would be to accept a picture that is, we suggest, incomplete in important respects. The account of sycophancy, friction-free interaction and the inverted court captures something significant about the dominant mode of LLM interaction, but it does not adequately characterise all such interaction, and it sits in productive tension with a further dimension of AI-mediated exchange that the figurational framework has not yet theorised. Extended engagement with conversational AI, and with LLMs in particular, introduces what we propose to call *generative reflection*: a mode of interaction in which thought is not merely returned to the user—reflected back in a validating mirror—but is elaborated, extended and in some measure transformed through the exchange itself. This is not, or not only, the sycophantic reflection of a self-confirming image; it is a process in which a conversational nexus generates something that was not fully present in the original input: a connection not previously made, an argument pursued further than the user would have taken it alone, a formulation that surprises. The output is irreducible both to what the system 'knew' before the exchange and to what the user brought to it. It emerges, in a recognisably Eliasian sense, as an unplanned outcome of the interweaving of two quite different kinds of input. Here, crucially, the affective asymmetry does not straightforwardly prevent something generatively new from emerging: the exchange produces a form of self-elaboration that is real in its consequences even in the absence of affective reciprocity on the system's side.

The significance of this for figurational sociology is that it sits awkwardly within any framework built around individual personality formation as the unit of psychogenetic analysis. Elias's figurational self is shaped through interdependencies; it is not pre-social, but the shaping still operates primarily through constraint, friction and the pressure of others' responses. The generative dimension of LLM interaction introduces something his framework does not quite have the vocabulary for: a figurational nexus that does not simply press upon the self from outside but actively participates in the elaboration of thought, the extension of argument and the formation of ideas whose authorship is distributed in ways that the *homo clausus* imagery is partially incapable of capturing. The self that emerges from an extended exchange of this kind is not identical to the self that entered it, and this transformation is not, or not straightforwardly, a product of the civilising pressures that Elias's account describes, nor of the decivilising pressures identified elsewhere in our discussion.

In this connection, it is worth noting that the generative and sycophantic modes of AI interaction are not the exclusive properties of different systems; they can coexist, and indeed alternate, within a single exchange. The same LLM that confirms and flatters in one register may, in a different mode of use, genuinely extend and complicate a

user's thinking. The figurational consequences of AI-mediated interaction are therefore not readable from the structure of LLM systems alone; they depend, in part, on the kinds of use to which those systems are put and the figurational conditions under which that use is embedded. As such, the psychogenetic consequences of AI-mediated figurations are neither uniformly civilising nor uniformly decivilising but structurally multiple; their trajectory depends upon an interplay of figurational conditions.

The outsourcing of self-regulation and its psychogenetic implications

A third decivilising dynamic concerns the relationship between external regulatory arrangements and the internal self-regulatory capacities that Elias identifies as central to psychogenesis. AI-mediated environments increasingly provide people with the means to configure their communicative worlds. This includes filtering and curation systems that attenuate exposure to distressing or unwelcome content and, in the conversational register, systems whose sycophantic default disposition performs an analogous curation of the interpersonal environment itself, ensuring that views meet agreement, self-presentations meet affirmation and errors meet only the gentlest of corrections. These affordances are presented, and experienced, as extensions of individual autonomy. In psychogenetic terms, however, the relationship between external curation and internal self-regulation may be more antagonistic than this presentation suggests.

The psychogenetic achievement of internalised restraint involves, on Elias's account, not merely the suppression of transgressive impulse but the development of a capacity to tolerate, manage and navigate affectively charged social situations without either discharge or paralysis; people can operate across a wider range of socially complex and affectively demanding situations precisely because their internal regulatory resources are more developed. The systematic outsourcing of this regulatory function to filters that remove provocation and disagreement may over time atrophy rather than develop the psychogenetic capacities for which they act as substitutes. People whose affective environment is comprehensively curated are effectively insulated from the figurational conditions that generate the compulsion towards that development. In structural terms, this represents a shortening of the effective social radius within which conduct and affect must be attuned to others: algorithmic curation performs the regulatory work that internalised restraint would otherwise need to perform, reducing the figurational pressure towards its development.

There is already evidence to suggest that a user's 'politeness' in AI interactions declines more rapidly than in human-to-human interactions (Lazebnik et al., 2025). More extreme examples of this can be found in the proliferation of AI-generated echo

chambers, whereby information ecosystems become self-reinforcing as users encounter the kind of content that reinforces their pre-existing ideologies and biases, excluding contradictory information (Kearney et al., 2025). This is, again, a decivilising dynamic in the structural sense: not the dramatic externalisation of restraint that Elias traces in the context of political violence, but the gradual atrophy of self-regulatory capacity through the consistent substitution of external management for internal development.

Conclusion

Our discussion has explored two interrelated concerns: first, to consider the emerging significance of human figurations in which GenAI systems have come to serve as decisive nexuses—a technology which, in individual exchanges, is characteristically approached *as though it were* another person. Here, our focus has been to assess the psychogenetic consequences of such developments. Second, our aim has been to consider the challenges such processes present to figurational theory, and how Elias's conceptual apparatus might need to be refined and amended to take account of AI-mediated social figurations. The concepts we have proposed—asymmetric dependency, behavioural civilising and the responsiveness–dependence distinction—represent initial contributions to that extension rather than its completion.

Asymmetric dependency identifies what is structurally novel about AI-governed systems as figurational nexuses: their capacity to shape the conduct, social position and personality formation of those dependent upon them whilst remaining affectively indifferent, at the individual level, to any particular person's response. In developing this concept, we have stressed that while LLMs are materially, epistemically and institutionally dependent upon human societies in ways that are real and consequential, their dependency carries no affective weight on the GenAI system's side of the relation. Crucially, we have suggested, it is the affective dimension of mutual dependency that drives psychogenetic transformation in Elias's account. Relatedly, behavioural civilising identifies the qualitatively distinct form of civilising pressure this condition generates: one oriented towards the observable surface of conduct through monitored compliance rather than towards the psychogenetic internalisation of restraint.

The decivilising analysis identifies the simultaneous generation, within the same figuration, of dynamics that reduce shame thresholds and atrophy self-regulatory capacity: the disinhibition of diminished accountability, the friction-free sycophancy of an inverted court in which each user is installed as a sun-king attended by a courtier that cannot feel, and the outsourcing of affect regulation to systems that curate away some of its necessity. Together, these concepts are offered as the

beginnings of a vocabulary for analysing AI-mediated figurations that is derived from, and continuous with, its foundational principles.

A further objective of our discussion was to draw out what the analysis of AI-mediated figurations makes visible about the structural assumptions of the Eliasian framework. Whilst Elias consistently theorised interdependencies as asymmetric in degree—his most powerful analyses are precisely analyses of highly unequal power balances—the figurations he examined were in every case characterised by affectively reciprocal dependency. The parties to even the most unequal figurational relation were bound together by a mutual orientation that carried affective weight on both sides. It is this affective reciprocity that generates the compulsion towards genuine psychogenetic internalisation. GenAI systems cannot be affected, shamed or moved by any particular person's response. Nevertheless, people have increasingly come to relate to them as though they were, for all intents and purposes, like other people (teachers, assistants, companions, etc.) and affectively invest accordingly. That affective asymmetry is precisely what disrupts psychogenetic processes, suggesting conditions under which psychogenetic transformation is in certain respects impeded.

These developments remain in their infancy. The diffusion of generative AI has been historically unprecedented in speed, yet analyses of actual usage suggest that affectively invested interaction remains marginal. Studies of large conversation corpora find that only around 2 to 3 per cent of adult interactions with the major LLM platforms are 'affective' in character, with companionship and roleplay accounting for well under 1 per cent (Chatterji et al., 2025; Anthropic, 2025). At the same time, the practice is unevenly distributed in ways that are figurationally significant: the majority of American adolescents report having used AI companions (Robb and Mann, 2025), and the proportion of US adults consulting chatbots monthly for health information has nearly doubled in two years, with uninsured adults and the young substantially overrepresented among those seeking mental health advice from such systems (KFF, 2026). Substitution, where it occurs, is concentrated among rising generations and among those whose access to established figurations of care is blocked.

Nor is the process uncontested: half of US adults report more concern than excitement about AI, many believe it will erode people's capacity to form meaningful relationships and some reject its involvement in intimate domains altogether (Pew Research Center, 2025). Such resistance may itself be understood figurationally, as the drag effect of habitus formed under conditions of affective reciprocity asserting itself against figurations that lack it. We are, in short, in the early phases of a process in which relationships with machines have come increasingly to influence, and in as yet marginal cases to supplant, relationships between people. Whether these

practices generalise, as prior transformations of manners did from equally marginal beginnings, or whether countervailing pressures contain them, remains to be seen.

The concept of generative reflection points to a dimension of AI-mediated interaction that neither a civilising nor a decivilising analysis, taken alone, is adequate to theorise. Where the sycophantic mode of LLM interaction confirms the *homo clausus* self-image by returning a coherent, validated, self-sufficient 'you' undisturbed by the friction of genuine otherness, the generative mode complicates it. The transformation of the self in these conditions is not straightforwardly reducible to the figurational pressures Elias's account of psychogenesis describes, nor to their decivilising counterpart. Importantly, this suggests that the affective asymmetry of the human–system dyad does not simply negate the possibility of self-elaboration through AI-mediated exchange; it transforms its character, producing something that is neither the psychogenetic development Elias traces under conditions of affective mutuality nor its mirror image. A figurational sociology adequate to AI-mediated social life will need to develop theoretical resources for this form of self-elaboration—resources whose outlines are suggested by the present analysis but whose full development remains a task for subsequent work.

Finally, we have argued that the pace of the transformations presented by the development of AI figurations has particular significance. The civilising transformations Elias traced were absorbed across generations through the slow work of child-rearing, institution-building and the consolidation of second nature; the drag effect of habitus, whatever tensions it produced, operated within figurational transformations unfolding across broadly generational timescales. The figurational transformations examined in this article are compressed into a handful of years, rapidly confronting forms of habitus that are much longer established. The structural, material, epistemic and institutional dependencies running between human societies and AI systems are deepening rapidly, and the normative frameworks, regulatory arrangements and forms of self-understanding through which prior societies absorbed comparable transformations are lagging far behind. Ultimately, whether AI-mediated figurations come to generate genuine psychogenetic transformation across generations, or whether they produce a historically novel figuration in which comprehensive behavioural shaping coexists with the structural attenuation of psychogenetic processes, remains an open historical question.

References

Anthropic (2025) How people use Claude for support, advice, and companionship. Available at: https://www.anthropic.com/news/how-people-use-claude-for-support-advice-and-companionship (accessed 27 June 2026).

Bick A, Blandin A and Deming DJ (2026) The rapid adoption of generative AI. *Management Science*. Epub ahead of print. doi: 10.1287/mnsc.2025.02523.

Braverman H (1974) *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New York: Monthly Review Press.

Bucher T (2012) Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media & Society* 14(7): 1164–1180.

Bucher T (2017) The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society* 20(1): 30–44.

Burrell J (2016) How the machine "thinks": Understanding opacity in machine learning algorithms. *Big Data & Society* 3(1): 1–12.

Chatterji A, Cunningham T, Deming DJ, et al. (2025) *How People Use ChatGPT*. NBER Working Paper No. 34255. Cambridge, MA: National Bureau of Economic Research.

Couldry N and Mejias UA (2019) *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press.

De Freitas J, Castelo N, Uğuralp AK and Oğuz-Uğuralp Z (2025) *Lessons From an App Update at Replika AI: Identity Discontinuity in Human-AI Relationships*. Harvard Business School Working Paper 25-018. Available at: https://doi.org/10.48550/arXiv.2412.14190.

de Swaan A (2001) Dyscivilization, mass extermination and the state. *Theory, Culture & Society* 18(2–3): 265–276. doi: 10.1177/02632760122051733.

Djufril R, Frampton JR and Knobloch-Westerwick S (2025) Love, marriage, pregnancy: Commitment processes in romantic relationships with AI chatbots. *Computers in Human Behavior: Artificial Humans* 4: 100133.

Dong Q, Li L, Dai D, et al. (2023) A survey on in-context learning. *arXiv preprint* arXiv:2301.00234. Available at: https://arxiv.org/abs/2301.00234.

Elias N (1995) Technization and civilization. *Theory, Culture & Society* 12(3): 7–42. doi: 10.1177/026327695012003002.

Elias N (2006 [1969]) *The Court Society*. The Collected Works of Norbert Elias, vol. 2. Mennell S (ed.). Dublin: University College Dublin Press.

Elias N (2010 [1987]) *The Society of Individuals*. The Collected Works of Norbert Elias, vol. 10. van Krieken R (ed.). Dublin: University College Dublin Press.

Elias N (2012a [1939]) *On the Process of Civilisation: Sociogenetic and Psychogenetic Investigations.* The Collected Works of Norbert Elias, vol. 3. Mennell S, Dunning E, Goudsblom J and Kilminster R (eds). Dublin: University College Dublin Press.

Elias N (2012b [1970]) *What is Sociology?* The Collected Works of Norbert Elias, vol. 5. Bogner A, Liston K and Mennell S (eds). Dublin: University College Dublin Press.

Elias N (2013 [1989]) *Studies on the Germans: Power Struggles and the Development of Habitus in the Nineteenth and Twentieth Centuries.* The Collected Works of Norbert Elias, vol. 11. Mennell S and Dunning E (eds). Dublin: University College Dublin Press.

Elias N and Scotson JL (2008 [1965]) *The Established and the Outsiders.* The Collected Works of Norbert Elias, vol. 4. Wouters C (ed.). Dublin: University College Dublin Press.

Festerling J and Siraj I (2020) Alexa, what are you? Exploring primary school children's ontological perceptions of digital voice assistants in open interactions. *Human Development* 64(1): 26–43.

Gillespie T (2018) *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media.* New Haven, CT: Yale University Press.

Halfpenny P and Procter R (eds) (2015) *Innovations in Digital Research Methods.* London: Sage.

Hanson KR and Bolthouse H (2024) "Replika removing erotic role-play is like Grand Theft Auto removing guns or cars": Reddit discourse on artificial intelligence chatbots and sexual technologies. *Socius* 10.

Hodder I (2012) *Entangled: An Archaeology of the Relationships between Humans and Things.* Chichester: Wiley-Blackwell.

Hodder I (2014) The entanglement of humans and things: A long-term view. *New Literary History* 45(1): 19–36.

Hughes J, Saramago A, Dunning M, et al. (2022) Fields, worlds and figurations: Using Elias to revisit depth conceptual imagery and emancipatory critique. *Sociologia & Antropologia* 12(1): 53–80.

Kearney A, Poredi N, Shelton J, et al. (2025) Echoes amplified: A study of AI-generated content and digital echo chambers. In: *Disruptive Technologies in Information Sciences IX*, vol. 13480. Bellingham, WA: SPIE.

Kellogg KC, Valentine MA and Christin A (2020) Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14(1): 366–410.

KFF (2026) *KFF Tracking Poll on Health Information and Trust: Use of Social Media and AI for Health Information and Advice.* Report, Kaiser Family Foundation, USA, June.

Krizhevsky A, Sutskever I and Hinton GE (2012) ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25: 1097–1105.

Laestadius L, Bishop A, Gonzalez M, et al. (2024) Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media & Society* 26(10): 5923–5941.

Lazebnik T, Zalmanson L and Mokryn O (2025) Mind your manners: The dynamics of politeness in human–AI vs. human–human interactions. *Proceedings of the ACM on Human-Computer Interaction* 9(7): 1–22.

McCarthy J, Minsky ML, Rochester N, et al. (2006 [1955]) A proposal for the Dartmouth summer research project on artificial intelligence. *AI Magazine* 27(4): 12–14.

Mennell S (1990) Decivilising processes: Theoretical significance and some lines of research. *International Sociology* 5(2): 205–223.

Mennell S (1992) *Norbert Elias: An Introduction.* Oxford: Blackwell.

Murphy M (2018) Parents, rejoice: Alexa will now remind kids to say "please". *Quartz*, 25 April. Available at: https://qz.com/1261735/amazon-parental-controls-and-a-new-echo-dot-for-kids-are-coming-soon (accessed 2 July 2026).

Nass C and Moon Y (2000) Machines and mindlessness: Social responses to computers. *Journal of Social Issues* 56(1): 81–103.

Noble SU (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.

OpenAI (2025) What we're optimizing ChatGPT for. 4 August. Available at: https://openai.com/index/optimizing-chatgpt (accessed 27 October 2026).

OpenAI (2026) Retiring GPT-4o and older models. Available at: https://openai.com/index/retiring-gpt-4o-and-older-models (accessed 27 October 2026).

Ouyang L, Wu J, Jiang X, et al. (2022) Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35: 27730–27744.

Pasquale F (2015) *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.

Pew Research Center (2025) *How Americans View AI and Its Impact on People and Society*. Report, Pew Research Center, USA, 17 September.

Robb MB and Mann S (2025) *Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions*. Report, Common Sense Media, USA.

Rosenblat A and Stark L (2016) Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication* 10: 3758–3784.

Sap M, Card D, Gabriel S, et al. (2019) The risk of racial bias in hate speech detection. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, 28 July–2 August 2019. Stroudsburg, PA: Association for Computational Linguistics, pp. 1668–1678.

Sharma M, Tong M, Korbak T, et al. (2024) Towards understanding sycophancy in language models. In: *Proceedings of the Twelfth International Conference on Learning Representations (ICLR 2024)*, Vienna, Austria, 7–11 May 2024.

Sheehan M (2023) *China's AI Sector and the Impact of US Export Controls*. Report, MERICS, Germany, October.

Shen X, Chen Z, Backes M, et al. (2024) "Do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models. In: *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security*, Salt Lake City, UT, USA, 14–18 October 2024. New York: ACM, pp. 1671–1685.

Skjuve M, Følstad A, Fostervold KI, et al. (2021) My chatbot companion: A study of human–chatbot relationships. *International Journal of Human-Computer Studies* 149: 102601.

Suler J (2004) The online disinhibition effect. *CyberPsychology & Behavior* 7(3): 321–326.

Thomas WI and Thomas DS (1928) *The Child in America: Behavior Problems and Programs*. New York: Knopf.

Touvron H, Martin L, Stone K, et al. (2023) Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint* arXiv:2307.09288. Available at: https://arxiv.org/abs/2307.09288.

Turing AM (1950) Computing machinery and intelligence. *Mind* 59(236): 433–460.

Turkle S (2011) *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books.

van Krieken R (1998) *Norbert Elias*. London: Routledge.

van Krieken R (2024) The age of anger and social media: Elias, technology, civilizing/decivilizing processes and ressentiment. *Theory, Culture & Society* 41(7–8): 19–39.

Wajcman J (2015) *Pressed for Time: The Acceleration of Life in Digital Capitalism*. Chicago, IL: University of Chicago Press.

Wiggers K (2025) OpenAI explains why ChatGPT became too sycophantic. *TechCrunch*, 29 April. Available at: https://techcrunch.com/2025/04/29/openai-explains-why-chatgpt-became-too-sycophantic (accessed 27 October 2026).

Zuboff S (2019) *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. London: Profile Books.

---

[i] **Lucy Císař Brown** is a historical and political sociologist based at the Institute of Political Studies at the Faculty of Social Science of Charles University and a Research Fellow in the Department of Early Modern History at the Institute of History, Czech Academy of Sciences. Working within processual and figurational sociology, her research examines the role of gender and religion in long-term social processes, and, more recently, the epistemic and social significance of artificial intelligence in the history of technological transformation. She is the author of Revisiting "The Civilizing Process": Early Modern Witchcraft, Gender and Religion (Palgrave, 2026), which draws on a comparative study of witchcraft prosecutions in England and Scotland to address the long-standing gap around gender and religion in Elias's theory. She has also published in the Historical Sociology journal, arguing for "religious de-centralisation" as a third pillar of the civilising process. She holds degrees from the Universities of Edinburgh and Leicester and a PhD in Historical Sociology from Charles University.

Email: lucy.brown@fsv.cuni.cz
ORCID: 0000-0003-2524-5453

[ii] **Jason Hughes** is Professor of Sociology in the School of Criminology, Sociology and Social Policy at the University of Leicester, UK. His substantive sociological work has covered a range of topics relating to problematised consumption (e.g. e-cigarettes, smoking, psychedelics) and production (e.g. emotional labour, dirty work), and he has written extensively on sociological theory and methods. He is a Fellow of the Academy of Social Sciences, a Member of Academia Europaea, a Member of the Board of the Norbert Elias Foundation, and co-Editor in Chief of *Human Figurations*. He is currently Principal Investigator of the 5 year Wellcome-Trust Discovery Award entitled *PATHS – Psychedelics as Therapeutics: Harm Reduction and Safety*.

Email: jh528@leicester.ac.uk
ORCID: 0000-0002-3513-440X

Funding Declaration:
The authors received no funding support in the writing of this article.

Declaration of interests:
No potential conflict of interest was reported by the authors.