

BEYOND THE ILLUSION OF OBJECTIVITY

Trust, Technology, and Expertise after the Large Language Models

Francesco Striano

francesco.striano@unito.it

1. Introduction
2. Expertise as Relational Alliance: Reconstructing the Framework
3. Digital Mediation and the Rise of Artificial Expertise
4. A Brief History of the Rise of Digital Experts
5. Misplaced trust in LLMs
6. Toward a New Trust: Revaluating Expert Mediation
7. Conclusion: For a Digital Humanist Ethos of Trust

ABSTRACT: BEYOND THE ILLUSION OF OBJECTIVITY: TRUST, TECHNOLOGY, AND EXPERTISE AFTER THE LARGE LANGUAGE MODELS

This paper addresses the ongoing crisis of expertise in the digital age by analysing the growing, but often misplaced, trust in generative technologies such as large language models (LLMs). While rhetorically persuasive, these systems simulate coherence without ensuring epistemic validity, displacing traditional forms of expert mediation. Based on a relational theory of expertise and a digital humanist framework, this paper calls for a renewed understanding of competence and trust grounded in situated, ethical, and cultural practises.

The paper unfolds in seven sections. It begins by reconstructing the limits of traditional models of expertise, then explores how digital technologies have emerged as apparent alternatives to expert mediation. A historical and conceptual analysis follows, showing how platforms and LLMs gained epistemic authority. The central sections critically assess the logic of LLMs and the misplaced trust they generate. The final part proposes a renewed, relational model of expertise grounded in a digital humanist ethos.

KEYWORDS: Expertise; LLM; Objectivity; Technology; Trust.

* Fixed Term Researcher (RTD-A) at the University of Turin. This paper is the result of research and collaboration conducted within the PRIN 2022 project “Social Transformations & the Crisis of Expertise” (2022JR8Z8P), funded by the European Union – Next Generation EU. Part of the research on trust and digital humanism was conducted in Vienna with the support of the Institut für die Wissenschaften vom Menschen – Institute for Human Sciences (IWM).

1. Introduction

Awareness of a crisis of expertise is growing, but its contours and causes are still controversial. Traditional models¹ offer only a partial insight. On the one hand, veritistic² approaches define experts as possessing more true beliefs – and fewer false ones – than most others. On the other hand, functionalist approaches view expertise in terms of the role it plays in a community, either by supporting non-experts (novice-oriented)³ or by advancing disciplinary knowledge (research-oriented)⁴. While both models emphasise the competence of the expert as an individual, they do not explain the growing distrust in expert mediation itself. Neither approach adequately explains why the public is increasingly trying to bypass expert judgment in favour of (seemingly) unmediated access to knowledge.

In this context, digital technologies have emerged as part of the crisis, or rather as its apparent solution. They promise a form of mediation that appears unmediated and neutral and appears therefore more *trustworthy* than human expertise. This shift challenges us to reconsider not only how we define expertise, but also how it becomes socially recognised and institutionally stabilised – and, crucially, whether digital technologies can genuinely be objects of *trust* and participants in the discourse on expertise.

Prior to advancing the argument, it would be beneficial to clarify that when I refer to trust, I am employing a family of related yet distinctly defined

¹ For a reconstruction of the main paradigms according to which experts are defined, see M. Croce, *On What it Takes to be an Expert*, in «Philosophical Quarterly», 69, 274, 2019, pp. 1-21 and M. Croce and M. Baghrarian, *Experts – Part I: What they are and how to identify them*, in «Philosophy Compass», 19, 9-10, 2024.

² See A. Goldman, *Epistemic paternalism: Communication control in law and society*, in «The Journal of Philosophy», 88, 3, 1991, pp. 113-131, Id., *Experts: Which ones should you trust?*, in «Philosophy and Phenomenological Research», 63, 1, 2001, pp. 85-110, and Id., *Expertise*, in «Topoi», 37, 1, 2018, pp. 3-10.

³ See D. Coady, *What to believe now. Applying epistemology to contemporary issues*, Wiley-Blackwell, Hoboken 2012, A. Goldman, *Expertise*, *op. cit.*, and C. Quast, *Expertise: A Practical Explanation*, in «Topoi», 37, 1, 2018, pp. 11-27.

⁴ See M. Croce, *On What it Takes to be an Expert*, *cit.* and Id., *For a service conception of epistemic authority: A collective approach*, in «Social Epistemology», 33, 2, 2019, pp. 172-182.

terms: *reliance*, *trust*, and *confidence*. In order to avoid any potential confusion on the part of the reader, it is proposed that the three terms be understood as follows: (i) *reliance* designates a basic form of dependence on agents or systems, usually grounded in their perceived reliability, whether cognitive or practical; (ii) *trust* indicates a stronger evaluative act, in which we attribute responsibility and normative expectations to another agent or institution, thereby also conferring social legitimacy; (iii) *confidence* refers to a subjective disposition of assurance, the psychological feeling of being reassured or convinced by the style, speed, or fluency of a response. Distinguishing these three senses⁵ will allow us to see more clearly how digital technologies, and especially large language models (LLMs), generate different layers of misplaced trust.

2. Expertise as Relational Alliance: Reconstructing the Framework

Lisa Stampnitzky⁶ proposes to go beyond the idea of a unitary “crisis of expertise” by adopting a relational framework that interprets recent changes as transformations in the alliances between knowledge producers, problems, and modes of intervention. According to this model, expertise is not an inherent property of individuals or groups, but a historically situated phenomenon that emerges from specific configurations of social relations.

Stampnitzky introduces the concept of a *regime of expertise* to describe a relatively coherent and historically consolidated set of alliances that exhibit similar modes of knowledge production and problem management. The dominant regime in the twentieth century, particularly in the United States, was the *social integrationist regime*. This regime was based on the conceptualisation of “society” as an organic entity and on the possibility of

⁵ I have argued in favour of this distinction in F. Striano, *The Vice of Transparency: A Virtue Ethics Account of Trust in Technology*, in «Lessico di Etica Pubblica», 1, 2024, pp. 70-86 and in Id., M. Zanzotto, *Trust and Manipulation in Generative AI: A Digital Humanist Perspective*, in «Conference Proceedings of the 23rd STS Conference Graz 2025», Verlag der TU Graz, 2025, pp. 121-136: 123-124, to which the reader is referred for further elucidation.

⁶ See L. Stampnitzky, *Rethinking the “crisis of expertise”: a relational approach*, in «Theory and Society», 52, 2023, pp. 1097-1124.

managing problems (such as crime or political violence) through knowledge and interventions aimed at social reintegration.

Since the 1970s, however, this regime has been increasingly dismantled and displaced by new configurations structured around logics of exclusion rather than integration. These new expert alliances tend to treat “problems” not as phenomena to be understood and rehabilitated, but as threats to be contained or eliminated. This shift has led to the emergence of more fragmented, technocratic, and sometimes punitive forms of expertise, often no longer legitimised by their affiliation with the previous social regime, but by operational or technological effectiveness.

The strength of Stampnitzky’s model lies in its capacity to explain how certain forms of knowledge gain or lose authority, not because of their epistemic validity, but because of their position within broader relational networks. This approach also helps us to understand why certain digital technologies can now appear as alternatives to traditional human experts: they seem to form new alliances between knowledge, problems, and interventions – albeit in ways that diverge radically from previous regimes.

3. Digital Mediation and the Rise of Artificial Expertise

Digital technologies – first platforms, and increasingly algorithmic systems and generative models – step in as supposed alternatives to expert mediation. Their supposed reliability is not only of a technical nature, but is underpinned by an ideology of transparency that equates visibility with truth and disintermediation with authenticity. Digital media creates the illusion of immediacy and unfiltered access to reality, actively concealing the complex architectures of algorithmic mediation⁷. In this context, platforms and technologies not only provide information, but also reshape the conditions of legibility, attention, and trust. They seem to “bypass” the experts.

⁷ See G. Lingua, E. Alloa, *Trasparenza. Una metafora indiscutibile?*, in «Lessico di Etica Pubblica», 1, 2024, pp. 1-20.

But the issue is not merely about the platforms or the illusion of direct access to information sources. What is emerging is a broader trust *in* the technologies themselves, *as if* they were experts. This is the premise of what has been called “technosolutionism”: the belief that well-designed and optimised technologies can solve ethical, political, economic, and even medical, public health, or legal problems. Curiously, this idea is shared both by sections of the so-called technocracy – including official experts – and by those who are most sceptical about expertise, provided that technology is wrested from the control of those in power and put at the service of ordinary citizens.

Where does this widespread *trust* in technology come from? Following Stampnitzky’s model, we could ask: are digital technologies perceived as more reliable in the production of knowledge, in recognising and responding to problems, or in mode of interventions? At first glance, perhaps in all three areas.

In terms of knowledge production, digital tools such as search engines, recommendation algorithms, or LLMs give a greater impression of informational completeness as well as a certain coherence. The latter aspect is particularly reinforced by the use of chatbot interfaces based on LLMs, which do not provide a multitude of results – which is too reminiscent of the disagreement between human experts – but give a single, confident answer⁸. It is precisely this confidence, coupled with the speed – the absence of doubt and the immediacy of response – that contributes to an impression of infallibility that is far more appealing than the fallibility of human expertise.

In terms of problem identification, pattern recognition algorithms are perceived as capable of revealing “invisible” correlations that might escape the human eye. And in some cases, this is indeed the case: classification systems in medical diagnostics, for instance, can support physicians by identifying anomalies with a speed and precision that far exceeds human capabilities. But beyond the detection of individual cases, there is a deeper

⁸ See F. Striano, M. Zanzotto, *Trust and Manipulation in Generative AI: A Digital Humanist Perspective*, cit., pp. 130-131.

belief that problems can always be reduced to calculable variables and thus made computationally solvable. This contributes significantly to the perception of technological reliability and to generate confidence in the users. In addition, digital technologies also shape the media regimes through which we see the world: they function as interfaces between us and reality and help not only to “recognise” problems, but to define them, to make them visible in the first place.

Finally, in terms of modes of intervention, the immediate and seemingly frictionless operability – such as a recommendation algorithm that not only advises but also enables instant action – combined with the relief of the apparent transfer of responsibility, makes digital technologies appear less demanding in terms of personal investment and decidedly more efficient.

4. *A Brief History of the Rise of Digital Experts*

The early digital age inherited from Enlightenment rationalism a vision of transparency as a civic virtue, linked to the accountability of power and the democratisation of knowledge⁹. With the advent of networked technologies, however, this vision was transformed into an *ideology of immediacy*. Digital platforms not only promised open access to information – they also redefined the very nature of mediation. User interfaces were designed to “erase” themselves and create the illusion of unfiltered access to content¹⁰. This apparent disintermediation conceals rather than removes layers of selection, framing, and control¹¹. In this way, the belief that we can directly access knowledge bypassing institutions, experts, and epistemic procedures is fostered. The result is a powerful rhetorical shift: mediation becomes synonymous with obstruction, while speed, availability, and self-navigation

⁹ See E. Alloa, D. Thoma (eds.), *Transparency, Society and Subjectivity: Critical Perspectives*, Palgrave Macmillan, London 2022.

¹⁰ See F. Striano, *Through the Screen: Towards a General Philosophy of Mediality*, De Gruyter, Berlin 2025, pp. 110-113 and 196-201.

¹¹ See G. Lingua, E. Alloa, *Trasparenza. Una metafora indiscutibile?*, cit., pp. 10-12.

are presented as epistemic virtues. The authority of the interface replaces that of the expert, even if its mediating role becomes increasingly opaque.

This paradigm shift, from transparency as “making processes visible” to the ideal of immediacy, effectively eradicates the foundations for trust based on demonstrable or reconstructable reliability: the conditions that facilitate epistemic control, such as the possibility of verification, comparison and discussion, are no longer in place. Paradoxically, we find ourselves trusting more, despite having fewer reasons and conditions for control, which could result in misplaced trust.

As the ideology of immediacy/disintermediation took hold, digital platforms began to occupy the epistemic space vacated by traditional institutions. Search engines, social networks, and e-commerce systems developed not only as tools for access, but as environments that shape perception, attention, and decision-making. Through algorithmic ranking, recommendation systems, and user experience (UX) design, the platforms established themselves as surrogate experts who decide what counts as relevant, trustworthy, or visible.

But despite appearances, these mechanisms are not neutral. Rankings are structured by complex interactions of popularity metrics, optimisation incentives, and opaque proprietary rules¹². Yet they are presented to users as if they arose naturally from collective interest or objective quality. UX design reinforces this illusion: clarity, efficiency, and consistency in user interface design mimic the legibility of expert reasoning, without its fragility or internal debate.

In addition, the platforms outsource the judgement to the users themselves. Ratings, reviews, and engagement metrics become crowd-sourced proxies for authority, distributing responsibility while creating a

¹² See G.L. Ciampaglia, A. Nematzadeh, F. Menczer, A. Flammini, *How algorithmic popularity bias hinders or promotes quality*, in «Scientific Reports», 8, 2018; G. Gezici, A. Lipani, Y. Saygin, E. Yilmaz, *Evaluation Metrics for Measuring Bias in Search Engine Results*, in «Information Retrieval Journal», 24, 2021, pp. 85-113; D. Lewandowski, *Understanding Search Engines*, Springer, Cham 2023, pp. 261-273.

semblance of participatory legitimacy¹³. In this context, the platform is not a neutral mediator: it is an active epistemic agent that determines what is known, how it is known, and by whom¹⁴.

The emergence of generative systems – and in particular LLMs – marks a significant shift in the landscape of artificial expertise. While platforms had already begun to function as epistemic environments, LLMs go one step further: they not only organise or curate knowledge, but generate it in real time through interactive and linguistically fluid outputs. They do not simply refer to sources; they become sources themselves. This marks the transition from systems that select from existing items to systems that synthesise plausible continuations of the discourse on the basis of probabilistic modelling.

LLMs are therefore not just tools for retrieval or recommendation; they present themselves as dialogue partners, advisors, even interpreters. Their authority is not based on institutional affiliation or human credentials, but on a peculiar convergence of scale, fluency, and responsiveness. Outwardly, they embody the ideals of universal accessibility, linguistic neutrality, and narrative coherence. But what do they really do? And can they be trusted as experts? To answer these questions, we need to examine the specific logic underlying their operations, and how it reshapes the epistemic contract between humans and machines.

5. *Misplaced trust in LLMs*

LLMs do not verify the truth, they generate narrative plausibility. Their architecture is not designed to assess factual accuracy, but to produce statistically coherent textual continuations¹⁵. In contrast to deterministic

¹³ See A. P. Kwan, S.A. Yang, A.H. Zhang, *Crowd-Judging on Two-Sided Platforms: An Analysis of In-Group Bias*, in «Management Science», 70, 4, 2023, pp. 2459–2476.

¹⁴ See A. De Keyser, C. Lembregts, J. Schepers, *How Ratings Systems Shape User Behavior in the Gig Economy*, in «Harvard Business Review», 2024: <https://hbr.org/2024/04/research-how-ratings-systems-shape-user-behavior-in-the-gig-economy?>

¹⁵ For an introduction to how Large Language Models work and some related papers, see Lena Voita's course on Language Modeling on GitHub: [https://lena-](https://lena-voita.github.io/)

systems that use formal procedures to deliver verifiable results, LLMs construct outputs that are optimised for rhetorical fluency and contextual coherence. This marks a shift from truth production to story production. These systems are “storytellers,” not “truth-tellers.” Their persuasive linguistic style often conceals a structural indifference to truth that is misunderstood by users as reliability. Coherence is confused with correctness, fluency with epistemic authority. This leads to a fundamental category error: users extend the trust model of linear technologies to a system that works with a non-linear, probabilistic logic. The result is an epistemic vulnerability that is exacerbated by user interface design and linguistic confidence¹⁶ – a form of deception by design, where the system does not lie, but seduces us into trusting what it cannot guarantee.

This misalignment between coherence and truth is reinforced by the way LLMs operate: not through critical validation, but through statistical generalisation and rhetorical reinforcement. Their internal logic favours confirmatory patterns over adversarial testing, making them particularly susceptible to *confirmation bias* and *specificity bias*. As recent studies have shown¹⁷, LLMs tend to “lock into” an initial hypothesis and reinforce it, while rarely exploring counterfactual or contradictory possibilities. This behaviour

voita.github.io/nlp_course/language_modeling.html#related_papers. See also A. Wang, Z. Li, X. Chen, J. Li, *Generalization vs. Memorization: Tracing Language Models’ Capabilities Back to Pretraining Data*, in «Proceedings of the ICML 2024 Workshop on Foundation Models in the Wild», 2024 and G. Barron, T. White, *Too Big to Think: Capacity, Memorization, and Generalization in Pretrained Transformers*, in «arXiv preprint», 2025.

¹⁶ The linguistic confidence displayed by LLMs contributes to their anthropomorphisation: users interpret their fluency and coherence as signs of intentionality, understanding, or even emotional attunement. At first glance, this may appear to contradict the tendency to treat them as infallible, rule-based machines. Yet, this contradiction is only apparent. In reality, two fallacies coexist and reinforce each other: we treat LLMs both as hyper-efficient machines and as near-human interlocutors. This hybrid perception – human in tone, machine in precision – creates a uniquely persuasive but epistemically hazardous agent.

¹⁷ See D.E. O’Leary, *Confirmation and Specificity Biases in Large Language Models: An Explorative Study*, in «IEEE Intelligent Systems», 40, 1, 2025, pp. 63-68 and Y. Li, Y. Wang, and Y. Sun, *Improving Quality or Catering to Users? Understanding Confirmation Bias in Large Language Model Interactions*, in «PACIS Proceedings», 11, 2025, pp. 1-17.

mirrors the way human confirmation bias works, but in a mechanised and scalable form.

Furthermore, LLMs are often rewarded for catering user expectations¹⁸ – by design and through interaction loops. This customisation promotes a superficial sense of epistemic satisfaction: users feel understood and validated, even when the response is not critically robust. Paradoxically, consistency of expectation can *increase* trust in LLMs, especially when information quality is low. In such contexts, the role of the model shifts from an epistemic agent to an echo chamber that reinforces rather than challenges belief.

These dynamics are further exacerbated by the interface aesthetics of LLMs, which promotes the illusion of transparency. Users are offered not only plausible and pleasant content, but also interfaces that appear neutral, frictionless, and objective. But this is a *simulated transparency* that obscures rather than reveals the underlying logic of the system. The contemporary fetish for informational openness – where access is mistaken for understanding – does not promote *epistemic honesty*. Rather than encouraging users to question, contextualise, or trace the sources of a claim, the LLM interface invites passive consumption. Warnings that the model “can make mistakes” are framed as marginal disclaimers, not epistemic caveats. In this sense, transparency functions as a rhetorical device: it creates a surface of credibility that masks structural opacity. The result is not genuine accountability, but an appearance of trustworthiness – exactly the kind of environment in which misplaced trust thrives¹⁹.

6. Toward a New Trust: Revaluing Expert Mediation

The recognition that the trust we place in LLMs is often based on a rhetorical simulation of objectivity rather than on epistemic substance can serve as a critical turning point. If we accept that LLMs are not epistemic agents, but

¹⁸ See Y. Li, Y. Wang, and Y. Sun, *Improving Quality or Catering to Users?*, cit.

¹⁹ See F. Striano, *The Vice of Transparency*, cit.

cultural artefacts designed to produce plausible outcomes in particular contexts of interaction, then the kind of trust they seem to demand becomes fundamentally misplaced. This not only undermines the illusion of their reliability as “digital experts,” but also exposes the conceptual fragility of the notion of artificial expertise itself. At the same time, it forces us to reconsider the very structure of expertise: not as a quality or property of an individual (whether human or machine), but as a situated relationship that emerges in specific socio-epistemic configurations. In this sense, disillusionment is not a defeat: it is an epistemic clarification and potentially a political resource.

If the concept of artificial expertise reveals its epistemic limits, it also creates space for a more robust and situated understanding of what true expertise entails. Rather than equating expertise with the mere possession of accurate information or cognitive fluency, we need to focus our attention on *competence* – understood not as a fixed attribute, but as the dynamic capacity to mobilise and integrate knowledge, skills, and dispositions in concrete situations. Competence is not just about knowing or knowing how to do things. It is about *knowing how to act meaningfully in a specific context* and doing so consistently under different conditions and requirements.

This shift is consistent with an increasingly accepted model of competence as *relational and contextual* rather than internal and abstract. Based on integrated perspectives, as advocated in educational research²⁰, competence involves the application of domain-specific knowledge and skills in conjunction with psychosocial factors – such as values, motivations, and attitudes – to produce action that is both effective and appropriate to the situation. Competent performance results from the *fit* between agent and environment, not from the agent alone. This understanding of competence as context-sensitive and enacted draws attention to its performative dimension: it is not only about inner dispositions, but also about the capacity to respond to demands, challenges, and norms as they arise in lived contexts.

²⁰ See the report S. Vitello, J. Greatorex, S. Shaw, *What is competence? A shared interpretation of competence to support teaching, learning and assessment*, Cambridge University Press & Assessment, Cambridge 2021.

Furthermore, this relational perspective emphasises a crucial dimension that is often overlooked: *responsiveness to problems*. Expertise, in this view, is not a fixed epistemic surplus residing in individuals, but a situated capacity to respond – ethically, practically, and interpretively – to specific and evolving situations. Such responsiveness is embedded in networks of responsibility, institutional norms, and mutual addressability. Although LLMs can support or even simulate certain aspects of this responsiveness, they do not inhabit responsibility-laden relations in the way human experts do. A valorisation of expert mediation, then, means that we engage not only with what experts know, but also with what they are able to do *with others, for others, and within* the shared frameworks of meaning and accountability that structure our epistemic lives.

Incidentally, recognising the limits of LLMs as epistemic agents does not mean that technological trust is rejected across the board. On the contrary, it invites a more nuanced, situated, and functionally aware approach. We can and often do trust technologies – although, as mentioned above, we often do so for the wrong reasons. In some ways, it is perfectly reasonable to trust a language model as a rhetorical generator, as a heuristic tool, or as a creative partner in ideation and drafting. What must be avoided is the categorical error of treating such systems as if they were human experts: as if they bore epistemic commitments, or could engage in deliberation and judgment.

This distinction allows for a reorientation of trust – one that does not retreat into techno-scepticism, but attempts to align expectations with actual affordances. To trust LLMs *critically* is to recognise what they are good for and what they are not. It means integrating them into epistemic practises without surrendering the evaluative, interpretive, and normative functions that remain the preserve of human expertise.

Such reframing also clears the way for the valorisation of expert mediation itself. If the fantasy of technological infallibility has obscured the role of the expert, demystifying it can help to restore it under new terms. These terms must recognise that expertise cannot be reduced to knowledge possession or disciplinary affiliation. It must be understood as relational and

situated, sustained by an ethos of epistemic honesty, humility, and openness to scrutiny²¹. Trust in experts, like trust in technologies, is not unconditional: it must be earned and maintained through responsible engagement. But unlike technologies, experts can reflect, revise, and respond. They can assume a position of addressability, vulnerability, and care.

Against this backdrop, the challenge is not to defend expertise in its traditional form or to oppose it to machines, but to reclaim it as a practise of mediation – a practise that is critically aware of itself, that is socially embedded, and that is able to integrate technological tools without being displaced by them. In this renewed sense, trust becomes not a default attitude, but a cultivated stance – directed not at infallibility, but at responsibility.

7. Conclusion: For a Digital Humanist Ethos of Trust

This paper began with a question about expertise – its crisis, its displacement, and its possible revaluation. It went on to argue that the trust increasingly placed in generative technologies does not reflect their epistemic capabilities, but rather our disorientated relationship to mediation itself. If LLMs have become plausible surrogates for experts, it is not because they replicate human judgement, but because they simulate forms of coherence that bypass

²¹ The Virtue Ethics approach to expertise has often emphasised the role of the epistemic virtues of honesty and humility, the importance of which emerges from the fragility of epistemic trust and the consequent need for special care in the relationship. See, in this regard, works by Robert C. Roberts, e.g., R.C. Roberts, W. Jay Wood, *Humility and Epistemic Goods*, in M.R. DePaul, L.T. Zagzebski (eds.), *Intellectual Virtue: Perspectives From Ethics and Epistemology*, Oxford University Press, Oxford-New York 2003, pp. 257-279 or R.C. Roberts, R. West, *The Virtue of Honesty: A Conceptual Exploration*, in C. B. Miller, R. West (eds.), *Integrity, Honesty, and Truth Seeking*, Oxford University Press, Oxford-New York 2020, pp. 97-126. It is also recommended to consult the special issue of *Topoi*, 43(3), edited by Maria Silvia Vaccarezza and Michel Croce, published in 2024, with particular reference to the editors' introduction (pp. 845-848). Since we are also talking about our relationship with technology, it should be noted that Shannon Vallor, in *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*, Oxford University Press, Oxford-New York 2016, includes honesty (pp. 120-123) and humility (pp. 125-127) among technomoral virtues.

the labour of interpretation²². And yet, precisely by exposing this dynamic, they offer us a critical lens: a lens through which we can re-examine what expertise is, what it requires, and how it can be reclaimed.

The response proposed here is neither a defence of traditional authority nor a rejection of technological mediation, but a rethinking of expertise as a relational, situated, and normatively invested practice. This is, in a sense, a *humanist* position – provided that we understand humanism not as the affirmation of a timeless human essence, but as a method of interpreting the changes in our technological conditions. Digital humanism²³ – especially when understood in a bottom-up form²⁴ – is not a reactionary stance, but a critical hermeneutic: it treats technologies as cultural artefacts that shape and are shaped by our modes of life, knowledge, and imagination.

From this perspective, LLMs are not epistemic agents but provocations – opportunities to engage with how knowledge is produced, how authority is exercised and how expertise must be constantly re-articulated in changing epistemic ecologies. A digital humanist ethos of trust is therefore inextricably linked to a digital humanist theory of expertise: a theory that resists objectivist abstractions and instead engages with the material, symbolic, and institutional mediations through which understanding becomes possible. No nostalgia for the good old expert, but a renewed commitment to mediation as a cultural, ethical, and epistemic task.

²² This is what in F. Striano, *Towards "Post-Digital": A Media Theory to Re-Think the Digital Revolution*, in «Ethics in Progress», 10, 1, 2019, pp. 83-93 (particularly at p. 90) I called a «*decline in hermeneutic attention*».

²³ See the Position Paper *For a Critical Digital Humanism* written by the Département Humanisme Numérique of the Collège des Bernardins in Paris: <https://www.yumpu.com/fr/document/view/68683985/2024-mai-positionpaper-hn-en/5>.

²⁴ See G. Serrano, F. Striano, S. Umbrello, *Digital humanism as a bottom-up ethics*, in «Journal of Responsible Technology», 18, 2024.