Quaderni di Teoria Sociale 1.2026
ISSN (online) 2724-0991 | DOI: 10.57611/qts.v5i1.646
www.riviste.morlacchilibri.com/index.php/quaderniditeoriasociale

Philip Di Salvo

Artificial Intelligence Inherits a Past. A Conversation with Gabriele Balbi

Abstract: In this interview, Prof. Gabriele Balbi reflects on the historical and material conditions of contemporary artificial intelligence (AI). Drawing on his work as a media historian, the conversation resists the prevailing tendency to apprehend the present moment in AI as historically unprecedented and instead situates the current "summer" of generative AI within the longer trajectory of digitalization, the platform economy, and the cycles of hype through which "revolutionary" technologies have repeatedly been announced and absorbed. From this perspective, Balbi argues that generative AI is less a rupture than a continuation: heir to the financialization, extractivism, and ideological framings of earlier digital formations, it can know only what has already been written, even as it promises to author the future. The conversation considers the political economy of the major AI platforms, the materiality of computation, the geographies of innovation beyond the Sino-American duopoly, and the paradox of a technology that announces the future while reinstating the extractive logics of the nineteenth century.

Keywords: artificial intelligence, media history, platformization

---

There is a recurrent tendency, in moments of conspicuous technological change, to perceive the present as historically unprecedented. The speed of transformation, the rhetorical force deployed by those propelling it, and the apparent reach of the technologies themselves combine to suggest that what is happening *now* bears little relation to what has come before. Yet, historians of media and technology have argued, with considerable consistency, that even the most consequential shifts rarely occur abruptly, and almost never without continuities with earlier formations (Gitelman, 2006; Balbi & Magaudda, 2018 – among others). The present moment in artificial intelligence offers a particularly clear illustration of this phenomenon. Generative AI has, in the space of barely four years, moved from a specialist concern to an everyday technology, and the velocity of its diffusion can produce the appearance of a substantive rupture. On closer inspection, however, the technology is deeply embedded in longer and more familiar discursive forma-

tions: in the trajectory of digitalization, in the consolidation of the platform economy, in the political economy of digital capitalism (Srnicek, 2025), and in the rhetorical apparatus through which "revolutionary" technologies are repeatedly announced and subsequently absorbed. There are, certainly, elements that are genuinely without precedent – most conspicuously the epistemic authority of systems capable of generating plausible knowledge claims at scale – but the significance of these novelties is most clearly legible against the continuities they unsettle.

Few scholars are better positioned to undertake such a reading than Prof. Gabriele Balbi, whose work has long been concerned with what is durable, recursive, and historically conditioned in the development of media technologies. Based at the Università della Svizzera italiana (USI) in Lugano, Switzerland – where he served as Rector ad interim in 2026 – Balbi has chaired the Communication History Division of the International Communication Association (ICA), having previously chaired the Communication History Section of the European Communication Research and Education Association (ECREA). His research traverses the history of telecommunications and digital media and the *longue durée* of communicative practice and infrastructural maintenance. His most recent monograph in English, *The Digital Revolution: A Short History of an Ideology* (Balbi, 2023), engages directly with the idea that frames this introduction. The book traces the genealogy of one of the most influential ideologies of recent decades: the claim that digitalization constitutes a revolution, a rupture with the past, a radical transformation of the human condition, and instead retraces where and how this conviction originated, which earlier revolutions it has consciously or unconsciously drawn upon, what narratives it has propagated, and which of its key elements have shifted while others have persisted and recurred across distinct historical periods.

This concern with continuity is central in Prof. Balbi's work, and the conversation that follows draws on a longer line of inquiry. In 2018, he edited a monographic section of *Quaderni di Teoria Sociale* titled "La svolta apocalittica" negli studi sul digitale, in which Balbi argued that scholarly engagement with digital media had shifted decisively away from the celebratory register of the 1990s and early 2000s turning instead towards what he termed, after Umberto Eco (1964), an "apocalyptic" turn (Balbi, 2018). The introduction reconstructed a canon of

now "classics" of digital critique – including Evgeny Morozov on solutionism (2013) and Christian Fuchs on platform exploitation (2014) – and argued that, although the section did not endorse the apocalyptic position wholesale, the critical orientation had become the prevailing mode of scholarly engagement with the digital. At the time, I contributed to the special issue with a reflection on how Edward Snowden's 2013 revelations represented a turning point in society's relationship to, and sense-making of, anything digital (Di Salvo, 2018).

This interview takes up several of the lines of inquiry developed in 2018 and follows them into the territory of generative AI. Yet, the conversation also surfaces a paradox specific to this latest technological formation, as generative AI sustains a peculiar relationship with the past (Kidd & Nieto McAvoy, 2023) that is taking shape in an already nostalgia-filled cultural and technological climate (Gandini, 2020). Trained on vast archives of texts, images, and cultural artefacts produced over decades or centuries, these systems generate outputs that purport to address the present and project the future, while in fact recombining materials drawn entirely from what has already been written, drawn, or recorded. The future they appear to inaugurate is thus assembled from the past – a paradox that throws into sharper relief the broader argument that runs through Balbi's work, namely that there is no clean rupture between what is and what was, and that the apparently new is invariably a reorganization of inherited materials. What this conversation adds in 2026 is also the recognition that the apocalyptic register has by now migrated from academic into ordinary public discourse, where it has solidified into a broader cultural climate across the Western public sphere.

Within an environment shaped by hype – in both its enthusiastic and its pessimistic inflections (Bender & Hanna, 2025; Galanos & Stewart, 2024) – *doomerism*, the conviction that artificial intelligence portends civilizational catastrophe and perhaps human extinction, has now come to occupy a conspicuous position in public and mainstream debates, often articulated through a corporate vocabulary that lapses into sensationalism precisely where it claims to address the so-called "existential risks" of AI. The temptation of these sublime registers, whether redemptive or catastrophist, is one that critical scholarship ought to resist. What is required, rather, is a sustained effort to cultivate measured, evidence-based, and historically grounded debate around artificial intelligence and its political-eco-

nomic foundations and its *real* and non-speculative harms, and to potentially acknowledge, as Thomas Dekeyser has argued in *Techno-Negative* (2026), that every technology has always emerged in tandem with its own negation. The interview took place over the Web, in Spring 2026.

I would like to ask you to characterize the historical moment in which we currently find ourselves within the history of artificial intelligence. Scholars typically describe its development in terms of "winters" and "summers": how, then, are we to understand the present phase?

Let us reason historically. The notion of artificial intelligence emerged in the late 1940s and, more decisively, in the 1950s, in part as a term defined in opposition to cybernetics, which was then the dominant framework and keyword; indeed, the name AI was introduced precisely to distinguish the new field from cybernetics. Already in the 1950s and 1960s, two paradigms came to be juxtaposed: *symbolism* and *connectionism*, as we may provisionally label them. The first prescribes that all instructions be supplied to the machine from above through systems of explicit rules: the worlds in question are simplified and self-contained, small environments in which artificial intelligence, by virtue of pre-given instructions, is required to produce a desired output. For a considerable period, this constituted the dominant paradigm. The second, which is now central, also originated in the 1960s and corresponds to machine learning and to its intersection with neural networks: machines are required to learn from the experience they accumulate, and neural networks are designed to simulate such learning processes. The machine, in other words, must extract patterns and regularities from data through a form of learning that is not necessarily supervised. These two paradigms have proceeded in parallel throughout the history of artificial intelligence; large language models, in this respect, do not constitute a novelty but rather inscribe themselves within this second paradigm.

We are therefore witnessing a marked resurgence of this second paradigm and, as you suggested, the return of one of the most luminous summers in the history of artificial intelligence. Scholars such as Simone Natale and Andrea Ballatore (2020), and many others remind us that AI moves through summers and winters: phases in which it is intensely discussed and presumed capable of transform-

ing society, and phases – one need only consider the 1970s, 1980s, and in part the 1990s – in which it nearly disappears from public discourse, ceding the centre of the stage to other keywords and concerns. We are unquestionably in a new summer. Many take the launch of ChatGPT in 2022 as its symbolic inception; in fact, however, the resurgence had already begun in the 2010s, with episodes such as the defeat of the world Go champion by Google DeepMind's AlphaGo (Bory, 2019). From 2021–2022 onwards we have seen Midjourney, ChatGPT, Gemini, and the various other applications. These products and names already feel familiar, even dated, yet they are in fact extremely recent. All have been bound up with the return of the second paradigm, in which machines, through the relentless processing of data, generate probabilistic patterns and produce credible, plausible, well-crafted, and well-written texts and images that appear strikingly human, perhaps excessively so, almost as a hyper-simulacrum.

I tend to view generative artificial intelligence in continuity with the history of the web and of the platforms. Do you regard this as something firmly embedded in our collective imaginary of AI, or is there still a tendency to conceive of AI as something that emerges outside those genealogies?

In the prevailing imaginary, in my view, it emerges from outside: these companies appear to us as supremely innovative, severed from any ties to the past. In reality, several forms of continuity are at work. First, it is precisely the firms of the "older" digital era – the so-called 1.0 phase – namely Microsoft, Amazon, Apple, and Google, among others, that invest most heavily in artificial intelligence. Companies established in the 1970s, such as Microsoft and Apple, or in the 1990s, such as Google, have unquestionably developed and continue to invest enormously in AI, even though it is becoming increasingly evident, especially in recent months, that the sector remains comparatively unprofitable. Second, generative systems could not exist without the web: it is precisely the proliferation of available data that enables these machines to process the vast quantities of information on which they depend. Without training conducted on the corpus of accessible web pages, none of this would have been possible. Other factors must also be considered: the growth of computational power – today's smartphones possess the capacity of supercomputers from only a few years ago –

and our improved understanding of the functioning of the human brain. A series of intertwined elements is therefore at play; yet without the World Wide Web and without the increase in the volume of available data, this form of generative intelligence (Fomasi et al., 2023). I believe, be inconceivable.

And perhaps the fact that AI is so deeply rooted in those genealogies has carried over into it a number of issues already characteristic of Web 1.0 and of its commercialization: digital power, centralisation, the entire question of data justice. Would you agree?

Yes, it seems to me that there are no major differences nor profound ruptures. We remain in an age of critique of digitalization: artificial intelligence is, fundamentally, part of this broader wave, even if it presents itself as a further caesura. If we wish to discuss the question of "revolutions," these are always cast as *the next big thing*: from this perspective, AI is undoubtedly the new great technology, just as in earlier moments we have had the revolutions of the computer, the internet, or mobile telephony, or, on a smaller scale, the CD-ROM, or the Metaverse, of which a great deal was said only a few years ago and which is now barely mentioned. There is, then, a constant dimension: AI is the fresh and innovative technology of which the digital revolution, however we wish to designate it, stands in need in order to remain culturally current. On the other hand, there is perhaps something further. Artificial intelligence has indeed succeeded in penetrating deeply into the everyday practices of a far broader population: notably in education, a sector with which we are actively engaged, but also in numerous other domains.

Almost every sphere of professional activity has been affected by AI mechanisms, even though several studies indicate that not all firms apply these technologies in the same manner, and that they do not necessarily prove economically advantageous: some research is even beginning to point to costs higher than those incurred by employing a corresponding number of human workers to perform analogous tasks. It is nevertheless true that one defining characteristic of AI is the breadth of fields it has affected: domains of knowledge, labour, everyday life, and entertainment. Nor should we overlook the extent to which it has shaped the social networks and social media: a substantial proportion of the content

circulating on these platforms is generated artificially and consumed by human users, which constitutes a further dimension of considerable interest. Recurring mechanisms are therefore at work, and your remark is unquestionably accurate: financialization is a fact that must never be lost from view. Far behind us is the moment of the second half of the 1990s, when the digital sphere was conceived as a new frontier and seemed destined to inaugurate new horizons of democratization: it is now apparent to all, even within mainstream and ordinary discourse, that we are dealing with corporations whose purpose is the generation of profit, and that this profit is realised principally on the basis of personal data. AI generates profit from data; yet a viable business model has not yet been identified.

I read a few days ago a striking figure on the very low percentage of users who subscribe to the paid versions of ChatGPT or comparable systems. The majority of users still rely on the free version; few subscribe to the paid tiers. If the subscription model fails to function, what remains? The familiar mechanism of advertising. This too is a continuity: it is no coincidence that we are now reading reports of these systems beginning to incorporate advertising into their architectures. The pattern is well established: if I do not pay for the product, if I do not purchase the newspaper, if I do not subscribe, what remains is advertising and my data.

It is, however, a curious circumstance: advertising embedded within the very machine that produces knowledge represents an additional order of concern. Encountering advertising while doomscrolling through Reels on Instagram or TikTok is one thing; but if the machine that generates content of a certain epistemological weight inserts advertising into that content, we are operating on a different plane altogether.

Yes, precisely because we are dealing with machines that, in many cases, simulate human interaction. It would be as if, in conversation with a friend, I was myself to insert advertising into the exchange. This was dramatized, you may recall, in *The Truman Show*: at a certain point the woman with whom Truman had shared his life begins to deliver advertising spots in his presence. Reading certain articles, I was reminded precisely of that film, a remarkable work of the 1990s, distinctly *televisual*, yet possessed of striking insights. The situation here is, of course, different. Nor should we underestimate the dimension of conver-

sational interaction: we are dealing with systems with which one may engage, for example, through WhatsApp; the continuous need for interaction is indeed met by these technologies. This represents a transformation of a particular kind: a transition from the imaginary friend to an interlocutor that is not embodied in flesh and blood but resides in circuits.

Speaking now from the perspective of a scholar of journalism, there is the long-standing question of the weight of hype (Lewis et al., 2025), of the "sublime" visions, in Vincent Mosco's sense (2005), both positive and negative, surrounding AI, often impelled by industry itself. There is, alongside this, the question of more direct harms: social impact, discrimination, and so forth. How do you assess the present cultural climate, considering the broader media discourse?

The question is a difficult one. It seems to me that mainstream media and, indeed, academia itself, have largely been towed along by the grand narratives of the major corporations. We are all turning to the study of AI, and this is evident within academia, historically slow to absorb novelty, where today numerous projects, monographs, and courses bear AI in their titles. We have followed this narrative: this technology will transform everything, and we can only study it, place it on the front pages of our newspapers, and disseminate it across mainstream media more generally. The major companies and platforms unquestionably possess the capacity to dominate public and academic discourses.

At the same time, particularly in recent years, mainstream media, public discourse, and the universities have begun to articulate a vigorous critique of these corporations. On the one hand, then, we have followed their lead in shaping the terms of our reflection and, in my view, this constitutes ideology in the proper sense, a worldview transmitted to us, a frame that appears to me totalizing; on the other hand, there is increasing space, including within public opinion, for the critical interrogation of these dynamics. The critique may at times manifest itself in rather elementary forms – adagios such as "I receive too much advertising," "they are listening to me," and the like – while in more sustained registers some are turning their attention to the forms of exploitation that the digital sphere, and AI within it, is exercising upon individuals: I refer to data inequality and to all that we now know concerning the use of human labour to sustain the opera-

tion of the machine. Whether this discourse has fully entered ordinary language and the shared imaginary, I cannot say.

There is perhaps a certain weariness, which concerns not only AI but social media more generally: certain indicators point to a gradual, slow, but significant disengagement on the part of younger users. According to the most recent data I have read, this is attributable to the fact that the social platforms have become saturated with hostility, fake profiles, and false information, such that they no longer constitute spaces in which users feel as comfortable as they once did. This phenomenon does not relate directly to AI, or only partially: a great deal of artificially generated content is produced precisely through AI. It is nevertheless, I believe, a discernible trend, one that, I am curious to know whether you, too, have observed.

I would like to ask you something concerning the entrepreneurial figures of the field. If we consider the great CEOs, so to call them, Steve Jobs continues to function as the paradigm of the previous era: a near-canonized figure, in both senses of the term. Today, by contrast, the leading names in AI – Altman, Bezos, Musk, and others – are framed in markedly more critical terms, apprehended above all through the lens of their capitalist interests rather than their technological contributions. Do you see significant differences between the first wave of Silicon Valley entrepreneurs and the present one?

I would say there are both similarities and differences. Let us begin with the similarities: these figures emerge from the Californian ideology, and even were they not to emerge from California, the substantive point, in my view, is simply that business is the foremost concern. The objective is the construction of a new digital capitalism which, like every form of capitalism, will rest on a structure of exploitation on the one hand and accumulation on the other. This has nothing to do with any form of digital philanthropy: we should dispense with that notion at the outset. This, it seems to me, is the principal point of continuity. There is, however, an important difference, I believe, in the sheer scale of the population that uses these technologies. If we return to Steve Jobs, to the 1980s, or to the early 1990s, very few people made use of a computer; somewhat more began to use the iPhone; today these technologies are in everyone's pockets, accessible to all. The phenomenon is far more global in its reach than was the case in the past.

A second similarity, of a political nature, and here the matter is open to discussion, is that in recent years the digital-capitalist gurus appear to have engaged with political affairs to a far greater extent than was previously the case. If Silicon Valley possessed an aura that was almost "left-leaning" some decades ago, today its aura is unquestionably more "right-leaning." I am, however, persuaded that with the next political shift their political creed will shift accordingly: I do not believe they hold any deep principled investment in this domain, save, and this is perhaps the most striking aspect, a certain servility before political power. They possess the resources and the potential to exert political leadership not only over the United States but perhaps over the entire world, and yet they refrain from doing so, at least in any conventional political sense; they pursue this end through other means, and this is open to debate. What strikes me as a notable difference, in any case, is this desire for proximity to traditional politics. Trump offers a clear illustration. So too do Altman and others; one thinks of that gathering around the table with Donald Trump and the entire array of digital capitalists, both old and new, Zuckerberg among them.

My own view is that the mask has fallen, in the sense that this entire culture, which is in fact of the right, had attempted to cultivate an imaginary that was not of the left but rather more broadly (neo)liberal – of the Clinton or Obama variety – and as soon as the opportunity arose, with Trump, it revealed its true allegiance. On this question, Becca Lewis at MIT has written on the right-wing history of the ideology of Silicon Valley (Lewis 2025).

Indeed. And on the matter of the Californian ideology, it was, I believe, Barbrook and Cameron, or others writing in their vein, who observed that, in the 1970s, there was a fusion of *new left* and *new right* (Barbrook and Cameron 1996). This is a particularly interesting amalgamation: two extremes that, in one fashion or another, come into contact. Perhaps this aspect persists.

It occurs to me that we have spoken extensively of the United States and of Europe, but there remains the rest of the world. I would like to ask: to what extent are we failing to look elsewhere? That is, to what extent does our perception of AI, beyond China, which constitutes the second evident pole, take account of the Global South,

of Africa? Are we missing something significant? Are we perhaps committing the same error that was committed with previous technologies: by adopting the Californian frame, we may have lost sight of how things might have been profoundly different had they been imagined within non-Western cultures and imaginaries.

Today the dominant configuration is the China–United States contest: on one side, the Chinese model, which appears more efficient from the standpoint of energy consumption and of data storage, among other criteria. On digitalization more broadly, we are unquestionably in a moment of Sino-American confrontation. Historically, of course, the United States has been pre-eminent for prolonged periods, but at certain moments Japan, too, came to the fore – one thinks of the 1970s – as did the Soviet Union with respect to the development of computing, and so forth. One region of the world to which we have never directed our attention is Africa: one need only consider Nigeria, potentially among the most populous nations in the world within a few decades. With respect to artificial intelligence, we have never examined Africa; we have considered it in other technological or media-related terms, for example, in connection with cinematography or with the alternative uses of mobile telephony.

In all candor, at present, I have not encountered any account of capitals of the world in which an artificial intelligence is being conceived as an alternative to that of the two poles of China and the United States, with the possible exception of South Korea. The discourse appears to me flattened along this axis, with Europe scarcely figuring except in relative terms: Mistral notwithstanding, the French initiatives notwithstanding, we appear to be confronted with a duopoly even at the level of the imaginary. With a measure of national pride, given that we are based in Switzerland, one may note that the "neutral" large language model, as it has been designated, developed by the Swiss National Supercomputing Centre together with ETH and other partners, I refer here to the Apertus project, prompted the realization that there may exist another way. The exploitation of data, for example, need not be conducted to the detriment of the publishers from whom, as I noted earlier, books are "extracted" or copied. There may indeed be another way.

I find it fascinating that generative artificial intelligence has access only to the past: data drawn from the past are deployed to generate content that should, at least

potentially, concern the present or, more strikingly still, the future. This relationship between past and generation is one I find compelling. How do you assess it, with respect to what AI is in a position to know?

What is intelligence? And what, in some sense, is recombination? From what does the new derive? The questions are difficult to resolve. Some maintain that there exists an inventive genius that breaks definitively with the past; yet STS scholarship suggests that is not in fact the case: there is no "classical" form of invention, and invention is increasingly conceived as a recombination of the past. AI would, in this respect, be the perfect instance: it draws upon disparate elements, recombines them probabilistically, and thereby generates something new. The question must therefore be posed.

There is an argument which one frequently encounters, for example among architects: only yesterday evening some of them were discussing it with me. They contend: "We could never be replaced, because it is plainly evident when our students execute a project through artificial intelligence, given that it never produces anything genuinely new." What might constitute the "genuinely new" in architecture would, in any case, require careful definition. There persists this sense – whether it represents a defense on the part of older generations or something else, I cannot say – that if I were to compose a novel, AI could not write it as Italo Calvino would, although it might "imitate" Calvino, and so forth. What I find more compelling is the question of memory which you raised: these large systems have, after all, been trained on books, newspapers, and writings produced over decades or centuries, and have in some manner internalized the human and cognitive heritage. This is a rather important and revealing observation: without that heritage, without the memory that has been transmitted across generations, this artificial intelligence might not have come into existence. Had a different memory been transmitted, we would have an artificial intelligence of a different nature. Perhaps in other regions of the world, with different cultural traditions, different books, and different cultures, there exists another form of generative intelligence. I am thinking here of the written record, but the same applies to the image: the iconographic heritage differs appreciably across traditions, including in religious terms, a dimension that should not be underestimated.

The extractivism of AI, which has "ingested" the entirety of the internet without seeking copyright authorization - and far be it from me to defend copyright as such - now extends to shredded physical books as well as the Washington Post *reported in January (Schaffer et al., 2026): beyond the material consideration in itself, there lies a genuine policy of plunder applied to whatever is available. The notion of extractivism, the constant search for new reservoirs of information to exploit, here applies even to the physical world, no longer solely to digital data.*

And it is not the first time. Reading the article in question, it occurred to me that Google Books had already done something analogous. The notion of extractivism is to my mind a productive one: there is a material extractivism applied to books, an energetic extractivism – so much so that AI companies are envisaging the construction of nuclear facilities and are reflecting on small modular reactors in order to meet this voracious energy demand. Several interesting reports are available; some weeks ago, I came across a number of substantial analyses concerning the consumption of materials such as the so-called "rare earths". There is therefore an unequivocally material dimension to artificial intelligence, a dimension we tend to underestimate and which, according to the relevant studies, will begin to exert significant impact within only a few years: producing CO_2 emissions and considerable energy consumption. You will have noted that Trump has recently issued a directive to return to coal for the United States military sector (The White House, 2026).

A portion of this sector employs artificial intelligence: we shall therefore have an artificial intelligence powered by coal. The point is striking, for we are returned to the locomotive, to the first industrial revolution, in which coal was the central element. There is, then, a dimension of "primitivism" on the one hand, nineteenth-century practices, including the exploitation of persons, including extractivism, and on the other hand a high-handed dismissiveness in the name of hyper-novelty. This strikes me as the most interesting paradox, including from a historical standpoint: nothing new, in any case, Kate Crawford had already articulated as much in her book (2021), yet a significant paradox endures: the most innovative phenomena, the future as we imagine it, bring once more to light discourses and forms of dispossession that appeared to have been resolved two centuries ago.

References:

Balbi, G. 2018, *La "svolta apocalittica" negli studi sul digitale: introduzione*, «Quaderni di Teoria Sociale», 1, pp. 11-24.

Balbi, G. 2023, *The digital revolution: A short history of an ideology*, Oxford University Press, New York.Barbrook, R., Cameron, A. 1996, *The Californian ideology*, «Science as Culture», 6, 1, pp. 44-72.

Balbi, G., Magaudda, P. 2018, *A history of digital media: An intermedial and global perspective*, Routledge, New York and London.

Bender, E. M., & Hanna, A. (2025). *The AI con: How to fight Big Tech's hype and create the future we want*. Harper.

Bory, P. (2019). Deep new: The shifting narratives of artificial intelligence from Deep Blue to AlphaGo. *Convergence*, 25(4), 627-642.

Crawford, K. 2021, *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*, Yale University Press, New Haven (CT) and London.

Dekeyser, T. 2026, *Techno-negative: A long history of refusing the machine*, University of Minnesota Press, Minneapolis (MN).

Di Salvo, P. 2018, *Sorveglianza, hacking e crittografia. L'"effetto Snowden" e l'emersione del lato oscuro del digitale nella ricerca sui media*, «Quaderni di Teoria Sociale», 1, pp. 53-80.

Eco, U. 1964, *Apocalittici e integrati*, Bompiani, Milano.

Fomasi, M., Barcella, D., Benecchi, E., & Balbi, G. (2023). Genealogy of an archive. The birth, construction, and development of the World Wide Web collection at CERN. *Internet histories*, 7(3), 277-294.

Fuchs, C. 2014, *Digital labour and Karl Marx*, Routledge, New York and London.

Galanos, V., & Stewart, J. K. (2024). Navigating AI beyond hypes, horrors and hopes: Historical and contemporary perspectives. In A. Ponce Del Castillo (Ed.), *Artificial intelligence, labour and society* (pp. 27-46). European Trade Union Institute.

Gandini, A. 2020, *Zeitgeist nostalgia: On populism, work and the 'good life'*, Zer0 Books, Winchester (UK) and Washington (USA).

Gitelman, L. 2006, *Always already new: Media, history, and the data of culture*, MIT Press, Cambridge (MA).

Kidd, J., & Nieto McAvoy, E. (2023). Deep Nostalgia: Remediated memory, algorithmic nostalgia and technological ambivalence. *Convergence*, *29*(3), 620–640. https://doi.org/10.1177/13548565221149839

Lewis, B. 2025, *Headed for technofascism: the rightwing roots of Silicon Valley*, «The Guardian», 29 January, available at https://www.theguardian.com/technology/ng-interactive/2025/jan/29/silicon-valley-rightwing-technofascism.

Lewis, S. C., Zamith, R., Bunquin, J. B. A. 2025, *Technological hype and AI in journalism: Five functions and why they matter*, «Digital Journalism», pp. 1-12.

Morozov, E. 2013, *To save everything, click here: The folly of technological solutionism*, PublicAffairs, New York.

Mosco, V. 2005, *The digital sublime: Myth, power, and cyberspace*, MIT Press, Cambridge (MA).

Natale, S., Ballatore, A. 2020, *Imagining the thinking machine: Technological myths and the rise of artificial intelligence*, «Convergence», 26, 1, pp. 3-18.

Schaffer, A., Oremus, W., Tiku, N. 2026, *Anthropic 'destructively' scanned millions of books to build Claude*, «The Washington Post», 27 January, available at https://www.washingtonpost.com/technology/2026/01/27/anthropic-ai-scan-destroy-books/.

Srnicek, N. 2025, *Silicon empires: The fight for the future of AI*, Polity, Cambridge and Medford (MA).

The White House 2026, *Fact sheet: President Donald J. Trump strengthens United States national defense with America's beautiful clean coal power generation fleet*, 11 February, available at https://www.whitehouse.gov/fact-sheets/2026/02/fact-sheet-president-donald-j-trump-strengthens-united-states-national-defense-with-americas-beautiful-clean-coal-power-generation-fleet/.

Philip Di Salvo is a senior researcher and lecturer at the Institute for Media and Communications Management (MCM), University of St. Gallen (HSG), Switzerland. His primary research interests include investigative journalism, internet surveillance, the intersection of journalism and hacking, and black box technologies. Previously, he served as a Visiting Fellow in the Department of Media and Communications at the London School of Economics and Political Science (LSE) from 2021 to 2022. From 2012 to 2021, he held various research and teaching roles at the Institute of Media and Journalism, Università della Svizzera italiana (USI). In the summer of 2024, Philip was a Visiting Fellow at Harvard University. Between 2018 and 2020, he also taught as a Lecturer at NABA – New Academy of Fine Arts in Milan, Italy. Philip is the author of two books: "Leaks. Whistleblowing e hacking nell'età senza segreti" (LUISS University Press, Rome, 2019) and "Digital Whistleblowing Platforms in Journalism. Encrypting Leaks" (Palgrave Macmillan, London, 2020).