visual communication
ARTICLE

Semiotics of the black box: on the rhetorics of algorithmic images

MASSIMO LEONE
University of Turin, Italy

ABSTRACT

This article critically analyses the semiotic pathways through which the new aura of algorithmic images is constructed, an aura which stems not so much from what they represent or from how it is represented but from the halo of mystery surrounding the very productive genesis of such images. Even their creators, from their super-technological laboratories, claim that they cannot fully grasp their emergence from artificial intelligence. Analysing these statements in depth, as well as the attempts that these same laboratories conduct to 'unravel' the mystery of the algorithmic images that they themselves fabricate and disseminate, however, one is seized with the suspicion that this mystery and aura are not due to intrinsic technical causes, but rather to the particular socio-rhetorical context in which digital and technological frontier knowledge is produced today, especially in relation to artificial intelligence. The 'black box' so often evoked to translate the inexplicability of artificial intelligence visual products might therefore be nothing more than a rhetorical device to protect and enhance the real black box, that of productive and industrial secrecy. In this whole process of algorithmic construction of the aura, then, the rhetoric of the unknowable image intercepts and highjacks a very long-standing trend in human cultures, in which images are precisely delegated the semiotic task of circulating the sense of a mysterious, ungraspable and unfathomable meaning.

KEYWORDS

algorithmic esotericism • artificial intelligence (AI) • black box • digital images • reverse-engineering • rhetoric

> Why don't they make the whole plane out of that black box stuff? (Steven Alexander Wright; American stand-up comedian, actor, writer, and film producer)

Visual Communication 2024
Vol. 0(0) 1–26

DOI 10.1177/14703572241247120
Sage

INTRODUCTION

In the digital zeitgeist where algorithms dictate not just the operational but the perceptual dynamics of our interactions, this article embarks on a critical exploration of the dual nature of algorithms. At the forefront of our inquiry lies the intricate dance between their technical and rhetorical functions, a relationship that not only moulds but also manipulates the digital ecosystem. Central to our analysis is the concept of authority, a trait traditionally ascribed to human agency, now increasingly attributed to algorithmic processes. This shift prompts a reevaluation of how authority is constructed, perceived, and legitimized within the digital domain, particularly through the lens of visual rhetoric.

The advent of sophisticated algorithms has ushered in an era where visual representations are not just mere reflections of reality but are charged with the power to influence, persuade, and even dictate societal norms and values. The rhetoric imbued in these algorithmic visualizations carries a weight of authority, often going unquestioned due to the technical obscurity that shrouds its functioning. The shroud existed also before the booming of the AI, but now it is becoming thicker. This article posits that the authority now wielded by algorithms is a direct consequence of their creators' and communicators' ability to perform rhetorical acts through visual means, thereby engaging in what we term 'algorithmic persuasion'.

By dissecting the semiotic and rhetorical strategies employed in the construction of algorithmic images, we uncover a deliberate orchestration aimed at enhancing the aura and authority of these digital constructs. This process, often masked under the guise of technical necessity, reveals a deeper intent to influence and control the semiotic discourse. The 'black box' of algorithms, thus, is not merely a technical challenge but a rhetorical tool designed to mystify and elevate the status of these digital entities.

Through a comprehensive analysis that spans across the development and deployment of algorithms in a specific socio-technical context, we demonstrate how the rhetorical functioning of algorithms extends beyond mere communication to actively shaping perceptions and establishing new paradigms of authority. This dual functionality not only highlights the need for a more nuanced understanding of algorithmic operations but also calls for a critical examination of the ethical implications arising from the unchallenged authority granted to algorithmic systems.

As we navigate this complex landscape, this article aims to provide a framework for understanding the interplay between the technical prowess and rhetorical strategies of algorithms, urging a rethinking of the role of visual rhetoric in the digital age. Through this exploration, we seek to illuminate the ways in which algorithms transcend their coded boundaries, becoming authoritative actors in their own right, shaping not just digital but also social realities.

In the contemporary digital era, this article analyzes the complex relationship between algorithmic processes and the semiotics of images, focusing on how the production and interpretation of algorithmic images construct a new form of aura and authority. It scrutinizes the intricate mechanisms behind these images, not from their direct representation, but from the mystique that shrouds their creation in AI labs. This inquiry into the black boxes of artificial neural networks reveals that the fascination and mystery surrounding algorithmic images owe less to their intrinsic technical properties and more to a socio-rhetorical context that valorizes secrecy and innovation, particularly in the sphere of AI. The study dissects the dual role of such images: as technical marvels and rhetorical devices that redefine the boundaries of authority, authenticity and intellectual inquiry in the digital age.

RHETORIC AND TECHNIQUE

The adjective 'augmented' comes from the Latin *augmentum*, a derivative of the verb *augere*, 'to increase' (De Vaan, 2008: 61–62, sub voce *Augeo*). The root of this verb, *aug-*, which has also an equivalent in the Sanskrit *vaks-*, 'to grow', is found in many Latin lemmas indicating augmentation, often in a metaphorical sense, and in relation to activities that could be called 'semiotic'. It is found, for example, in the Latin word '*augur*', in English also 'augur', which is the one that causes the *omina*, the signs of the future, 'to grow' or 'to emerge' from the observed reality; but it is also detectable in *auctoritas*, and, thus, in *auctor*, that is, the one who, by producing signs, increases their own prestige.

This article intends to deal precisely with the relationship between *augmentum* and *auctoritas*, through some general semiotic considerations (Leone, 2013: Preface), but also in regard to a more circumscribed topic, namely the relationship between 'augmented reality' – and in particular 'the augmented image' – and the range of semiotic and epistemic positions related to 'sign augmentation', namely, authorship, authoriality, authority and authoritarianism (Leone 2013, *The Semiotics*). All four of these discursive regimes derive from operations of semiotic augmentation, but with very divergent relational outcomes.

This means not only that, to construct authorship, signs must be 'added' to the world, but also that these signs must possess a special character, an 'augmented' one, precisely. In what, for example, are augur signs 'augmented'? And in what ways do they help define the authorship, but also the authoritativeness and authority, of the augur? First, bringing out these signs from reality is not for everyone but requires technical expertise: where the ordinary Roman citizen simply sees a few birds in flight, the augur identifies what might be called 'patterns', that is, visual configurations, the result of the motor behaviour of a few birds in flight, which the augur indicates is a sign (Manetti, 1987). Today we know there is no causal relationship between these configurations and the future state of things (Leone, 2021). Certainly, these flights cause something,

for example they cause the augur's vaticinations, but they do not in any way eventuate in what these vaticinations represent, namely, the future. The operation of the augur, however, does not consist in identifying causes, as in the case of physicians, but rather in pointing out indexical signs. Through a technical ritual, the augur has the flight of birds in the sky, a natural occurrence, to be seen as an index of the future (Annus, 2010). The physician of ancient Rome also did the same, that is, he identified certain features of the sick person's body as indexical signs of their illness, except that these signs were also symptoms, that is, underlain by a physical–causal relationship (Cary, 2008). The augur did not identify new causal links in the world, for he did not claim that the flight of birds would cause the future, but he guaranteed, through his own technical ritual discourse, that that flight was a sign of the future, an effect caused by it.

Here, then, in essence, is the difference between the physician or scientist and the augur or author: the former identifies some signs as the effect of a past cause, while the latter identifies some signs as the effect of a future cause (Leone, 2015). It seems obvious to us today that anyone who identifies signs in the present caused by the future is a waffler. Signs can only be caused by the past: it is an axiom of scientific thought. There can be indexical signs of causes, not of effects. How, then, did the augur manage to pass off the flight of birds as a sign caused by the existence, in the future, of a not-yet-occurred state of affairs? The answer: by resorting to technical knowledge, as noted above, except that such knowledge was not like that of the physician, based on the observation of causes and effects in reality but, rather, a knowledge that could be called 'rhetorical'. If the physician's technique identifies in the present the signs caused by the past, the augur's rhetoric identifies in the present the signs caused by the future. This operation is accomplished precisely by presenting rhetoric as a technique, that is, as a doing motivated by the real, where it is, on the contrary, an arbitrary doing, in no way motivated by the real (Michaelstadter, 1913).

In fact, anyone could exercise the role of the augur, and the signs they identify could be identified in any other way and endowed with a completely different meaning. The strength of the augur's discourse consists in making people believe that this is not true; that the augur's field of action, in short, is not a rhetoric but a technique, which could not be exercised in any other way, and, especially, could not lead to alternative results. That rhetoric should pass itself off as a technique is patently absurd, yet the augur succeeds in it precisely because of the 'augmented' character, so to speak, of their performance: the augur's clothes, their gestures, their words, the whole rituality of their role and function have as their purpose to give the rhetoric of vaticination a technical value. Ritual is in fact a discursive regime that, by 'augmenting' the sign tenor of communication as compared to its non-ritual circulation, gives it an auratic character and, above all, attributes motivation to the arbitrary (Leone, 2011).[1]

But, take away the augur's toga, remove them from the altar, change their Latin to a vulgar dialect, do away with their gestures and intonation and, above all, strip them of the label of augur: once all these 'augmented' signs are removed, there will be nothing left of the rhetoric, hence the inability to present it as a technique: does that cross flight of a seagull indicate that there will be a storm tomorrow? And who says so? Who is the augur to claim that? Why that seagull and not that sparrow? And why a storm and not an earthquake?[2] On the contrary, take away the doctor's lab coat, tear off his stethoscope, deprive them of their social recognition as a doctor: when this doctor without any 'augmented' signs any more – this doctor without rhetoric – tells the patient that a certain spot on their skin is actually a basal-cell carcinoma, the latter will be the first to thank the former, not for their rhetoric, but for their technique.

The situation is actually more complicated, as in modernity it is now possible to construct authority not only by adding signs to the current discursive regime, but also by subtracting them. This is an effect of the third element at play besides technique and rhetoric, namely, criticism. As criticism debunks the rhetorical character of technique, that is, pseudo-technique, the latter reacts by developing minimalist rhetorics,[3] which work by subtraction rather than addition. But this is not the occasion to explore such subtleties. Suffice it to conclude, at the end of this first section, that semiotic augmentation is evidently relational with respect to the discursive community in which it occurs and involves a construction of both authorship and authority. The rhetoric of sign augmentation, however, is not motivated by the real, and, thus, owes its force to an implicit symbolic convention. This causes the signs of augurs of all time to be subject to decay and a kind of sign inflation. After a while, the augur's gestures no longer convince anyone; they are imitated and made the object of caricature and mockery; they cease to convince that, indeed, in the real there are signs of future states, *omina*. This decay is dramatically fast in modernity: signs that are used to lend prestige to a pseudo-technique lose their aura and must be quickly replaced, with an overall effect of devaluing the rhetoric of technique in general (and, sometimes, that of criticism with it).

Understanding how digital and algorithmic technologies gain recognized authority involves looking closely at two things: first, how symbols and meanings are enriched or expanded; and second, how authority is established through persuasive communication as well as practical methods. This detailed examination is essential to grasp how certain actions or messages, which are initially seen as persuasive attempts, eventually come to be seen as authoritative technical knowledge. Essentially, this approach underlines the importance of recognizing that technology's influence and legitimacy stem not just from its technical aspects but also from how it is communicated and perceived, merging technical skill with the art of persuasion.

This initial segment delves into the historical and etymological roots of 'augmentation' and its profound connection to 'authority', demonstrating

how augmentation in a semiotic sense – whether through the rituals of augurs or the discourses of modern technology – serves to construct and enhance authority. This process is intricately linked to the technical expertise and rhetorical strategies employed to make certain signs (e.g., the flight patterns of birds in ancient Rome) meaningful and authoritative. The distinction between the augur's and the physician's approach to signs – indexical versus symptomatic/causal – highlights the nuanced relationship between semiotics, authority and the perception of technical versus rhetorical knowledge.

DEVICES AND COMPUTATION

In late modernity, then, this inflation produces a paradoxical chiasm between rhetoric and technique: on the one hand, as it is increasingly difficult to construct a pseudo-technical enchantment of rhetoric, the sign augmentation required to achieve such a result must be ever greater, and, above all, it must increasingly simulate a technical language, which essentially relies on two discursive elements: on the one hand, the use of devices and, on the other hand, computation. Devices and computation are the means by which any rhetoric that produces indexes can be passed off as a technique that discovers causes. The reason is simple: since science often uses technique–that is, devices and computation–to discover causes in the world, then simply evoking these two elements is enough to give the impression that one is doing the same. The chiasmus has the dramatic consequence that these two elements, i.e., devices and computation, become subject to the same sign inflation that characterizes rhetoric: as pseudo-scientific claims appear that purport to arise from the application of devices and computation to reality, all scientific discourse begins to arouse suspicion (Lotman, 2019). In short, we have moved from a situation in which the augur was passing themself off as a doctor to a situation in which the doctor is being passed off as an augur.

All of this initial discourse on the epistemology of sign augmentation in relation to science and pseudo-science, rhetoric and pseudo-rhetoric, goes in a direction that the reader might have already guessed. Academic discourse, in particular that of the humanities and the social sciences, reflects a great deal on the 'augmented' character of signs, and especially of images, in numerous social spheres, from sentimental relationships to play; it also reflects on the extent to which these 'augmented images' intervene to corroborate the technical character of contemporary scientific discourse (Dondero and Fontanille, 2012); it is not always the case, however, that the critique of the humanities and social sciences turns in on itself to observe how often in their own discourse 'augmented signs' and 'augmented images' are resorted to precisely for the purpose of corroborating, if not constructing, the feeling of a technical dignity.

We are not implying that the use of big data, AI, digital imagery and data visualizations is merely a modern form of ancient augury, where predictions were made based on bird flights. Rather than casting a critical eye on these contemporary methods in the humanities and social sciences, we sug-

gest that they might not always directly correlate with real-world causality, like a doctor's stethoscope diagnoses symptoms. Instead, these technological tools often leverage the persuasive power of digital signs to create an impression of scientific accuracy around essentially interpretive analyses, where the signs used suggest causality without direct evidence.

In tackling the contemporary manifestation of the dynamic between augmentation and authority, particularly in the face of digital technology's evolution, this article argues that the inflation of signification – whereby devices and computational methods are employed to lend a veneer of technical authority to essentially rhetorical acts – mirrors the historical practices of augury but with modern tools. This inflation not only challenges the authenticity of scientific discourse but also blurs the lines between genuine technical innovation and rhetorical mimicry, raising questions about the legitimacy and authority of scientific claims in the digital age.

By framing the central hypothesis within the broader context of semiotic augmentation and authority, this article not only reaffirms the intertwined nature of rhetoric and technique in constructing authority but also critically examines how contemporary digital technologies, through their augmented visual and computational capabilities, contribute to this process. It underscores the need to critically assess the rhetorical underpinnings of technical claims, especially in an era where digital imagery and algorithms play a significant role in shaping perceptions of authority and legitimacy. This approach aligns with the introduction's emphasis on the dual role of algorithms as both technical tools and rhetorical devices, thereby setting the stage for a comprehensive exploration of their impact on authority and visual rhetoric in the digital age.

A CASE STUDY

On 17 June 2015, Alexander Mordvintsev and Mike Tyka (software engineers at Google) and Christopher Olah (an intern at the same company) published a post entitled 'Inceptionism: Going Deeper into Neural Networks' in the company blog dedicated to AI (Mordvintsev et al., 2015). As will be seen, this post was bound to be widely influential, but this should not distract the semiotician from pointing out the discursive framework in which the post was situated. First of all, Google AI Blog is neither an academic journal nor one of dissemination, but rather a publication tool of the Google company. Although this is self-evident, it is necessary to emphasize it in order not to overlook the fact that these corporate channels can indeed be used to disseminate new research, but this dissemination function is never separated from an advertising or marketing intent. It can be as intentional as it can be unintentional, in the sense that when something is communicated through these kinds of channels – even though they are not governed by classical academic procedures of scientific quality control (which often take far, far too much time compared to the corporate production rhythm) – it is nevertheless received as potentially interesting news for anyone concerned with the topic.

Already from the title, then, it is clear that the proposed communication is by no means aseptic, but instead taps its terms and concepts into a background heavily laden with metaphorical connotations. In this case, for example, 'Inceptionism' is a term derived from a work of fiction, namely Christopher Nolan's film *Inception* (2010), while the expression 'going deeper' evokes an exploration of the abysses; moreover, even the naming of these abysses to be explored, i.e. 'neural networks', is highly metaphorical and poorly motivated from a biological point of view: the so-called 'artificial neural networks' are indeed networks and they are artificial but they are very different from neural connections.

The content of the post cited above is well known to specialists: automatic image classification and automatic verbal recognition have made great strides; to accomplish these tasks, mathematical models have been created that seem to work, in the sense that they return increasingly satisfactory results. If one uses an automatic translator between two known languages, for example, one finds that its work is, most of the time, very satisfactory. The point, however, is that it is not entirely clear how this satisfactory result is achieved. The post is then essentially about some suggestions on the 'reverse-engineering' of such results, that is, an attempt at understanding, from them, how the neural networks that produced them worked.

One of the challenges of neural networks is understanding what exactly goes on at each of their layers. We know that, after training, each layer progressively extracts higher and higher-level features of the image until the final layer essentially makes a decision on what the image shows. For example, the first layer maybe looks for edges or corners. Intermediate layers interpret the basic features to look for overall shapes or components, like a door or a leaf. The final few layers assemble those into complete interpretations – these neurons activating in response to very complex things such as entire buildings or trees.[4] The post focuses, in particular, on image recognition: an artificial neural network is trained by exposing it to millions of examples and adjusting its parameters until it produces the desired results.

That a neural network produces the expected results, however, is not completely satisfactory, from a human perspective, because one also wants to understand how they were obtained, that is, what kind of artificial cognitive patterns have emerged from training the network. This is also crucial for one to be able to actually attribute human-like intelligence to the network, that is, to distinguish between a perfectly simulated intelligence and an intelligence that not only simulates the results of a natural intelligence but obtains them by following a criterion that is analogous to those followed by a natural intelligence. I will not dwell on this philosophical aspect, which obviously refers back to the concept of the Turing test (Santangelo and Leone, 2023).

The type of artificial neural network examined by the Google post typically consists of a millefeuille system in which there are 10 to 30 layers of networked neurons that progressively distill the final result (Wei et al., 2015). The system, again, seems to work, but it is not totally clear how the end result is

achieved. In other words, the artificial neural network produces the distillate, and one knows what the still that produces it looks like, but it is not easy to determine which drops go through it. From the way this millefeuille system responds during training, it seems that it proceeds through a kind of generative pathway, whereby the first, deeper layers of the network are trained to respond with respect to the presence or absence of what semiotics would call 'plastic formants', while the more superficial layers would become sensitive to what Greimassian semiotics would call 'figures', or lexicalizable objects (Greimas, 1984). In face recognition, for instance, the first layers see edges or blunt corners, while the last ones recognize eyes and mouths. A caveat must be introduced at this point, however: this way of describing the inner workings of an artificial neural network exploits an analogy with Greimas's generative path, but the analogy is obviously very imperfect; the path of generative semiotics, in fact, is not one of progressive gestalt refinement, but aims at accounting for the emergence of meaning; that is why it summons concepts such as enunciation or narrativity, which have no equivalent at all in the scheme of operation of artificial intelligence.

According to the post, then, one way to understand what happens in the transition between one layer of the artificial neural network and the next is to reverse the process and, thus, ask the neural network trained to recognize a certain type of image to start from an image full of visual noise and transform it until the neural network itself produces the desired effect. This is actually a well-known operation in art history, gestalt theory and perception theory, except that it has usually been accomplished in relation to human visual intelligence. Starting from a blank sheet of paper, or better yet from a chaos of unrelated dots, what strokes should one draw, or what dots should one join, so that the visual meaning of a face emerges? One dot is too little, but two dots surmounting a horizontal line are perhaps already enough. Merleau-Ponty had also conducted a similar thought experiment on the face: why is it that by reversing it we often find it monstrous (Bertrand, 2023)? Theoretically, Greimas's generative path could also be used for a 'reverse-engineering' operation; from a certain point of view, some literary experiments by writers close to Greimas's seminar, such as Calvino or Perec, have adopted such a perspective, that is, working craftily on the literary text, by small progressive adjustments, until the result indicated by the theory is reached. The exact Greimassian equivalent of this 'inceptionism' would be to ask a Greimassian semiotician who has never read the meticulous analysis of Maupassant's novella 'Deux Amis' – conducted by Greimas himself – to read the analysis without knowing what the novella is about, and to reconstruct it from the semiotic dissection of it.

The comparison between the operation of artificial neural networks (ANNs) and Greimas's generative path offers a fascinating exploration into the parallels and divergences between computational processes and semiotic theory. ANNs, structured in layers, progressively refine inputs (e.g. visual data) from basic elements to complex outputs (e.g. facial recognition), reminiscent of how Greimas describes the emergence of meaning from fundamental semi-

otic elements to complex narratives. However, this analogy highlights a crucial discrepancy: while ANNs process and identify patterns without an inherent understanding of meaning or context, Greimas' path is deeply rooted in the production of meaning, engaging with concepts such as enunciation and narrativity that have no direct counterparts in ANNs. This comparison further extends into the concept of 'inceptionism', where reversing the ANN's process to generate recognizable images from noise mirrors artistic and semiotic strategies of creating meaning from abstraction. Despite these parallels, the critical distinction lies in the ANN's lack of intrinsic interpretative capability, contrasting with the humanistic, meaning-centred focus of Greimas's semiotics, underscoring the intricate relationship and fundamental differences between machine-learning methodologies and humanistic theories of meaning.

The blog proposes producing progressive adjustments to the image so that the inverted artificial neural network recognizes in it those objects that it has been trained to recognize. One of the purposes of this procedure, according to the post, is that it can help identify and correct cognitive biases developed by the artificial neural network during its learning phase. The example proposed by the blog is often reported in the literature on the subject (Nguyen et al., 2016a, 2016b): if an artificial neural network was trained to recognize dumbbells from pictures that always showed them being lifted by muscular arms, the moment it is asked to produce one of these pictures using the same (inverted) mechanism by which it identifies them, it represents not dumbbells but dumbbells together with forearms. This reveals that, although the network's result may be satisfactory, because it does recognize a dumbbell in a picture every time it is asked to do so, in reality the network is not recognizing the dumbbell but rather a complex object composed of a dumbbell and a forearm, along lines of a visual intelligence of reality that follow mysterious internal rules.

This invisible discrepancy can however become visible in particular cases and manifest itself as an error that reveals the fact that, in reality, the network effectively identifies dumbbells in photos but does not understand what a dumbbell is. To be honest, even interaction between humans often occurs in this way.

An example: in French, the Italian word *olio* is translated as *huile*, except that the former is a masculine word, whereas the latter is a feminine one. Many Italians, often even after developing seemingly perfect French, do not know that *huile* is feminine, and they normally do not learn this from interactions with native speakers. An Italian native speaker can continue to hear the French pronounce *l'huile* and never suspect that it is a feminine word; at the same time, the Italian speaker can pronounce this expression correctly, while equivocating on the gender of the word that appears in it, and not give the impression that a mistake is being made, that is, that the speaker hosts in their internal intelligence of the French language the potential for an error. This in fact emerges as soon as *huile* is used with an adjective posited with the indefinite article: *j'ai acheté un huile essentiel* is a phrase that reveals the

fact that the speaker is not a native speaker and has never internalized that *huile* is a feminine word (the correct sentence would be: *J'ai acheté une huile essentielle*). Similarly, artificial neural networks can satisfactorily interact with humans when they have to recognize images, but when one asks them to produce images with the same criteria, then one realizes that their understanding of what identifies a certain object as such is satisfactory but partial, and therefore prone to error: a neural network that was shown a dumbbell raised with teeth, for example, might not recognize it because its visual and semantic segmentation of the dumbbell, unlike the human one, includes the hand that raises it, although this does not transpire in most recognitions, precisely because the dumbbell that prototypically appears in photographs appears there jointly with a hand and a forearm.

But what if, in this back-engineering process, one does not submit to the artificial neural network a chaotic image, but any image, which the network will then have to modify in order to transform it into the image that would ideally produce the identification for which the network was trained? It is here that the technical procedure for identifying how the layers of artificial neural networks work begins to give rise to a visual rhetoric. Indeed, since the first layers of the network react to plastic formants, if these layers are asked to modify any image, they tend to produce a version of it in which these same plastic formants are accentuated. (Figure 1) The effect is that of a transformation of the images in an Impressionist or Post-Impressionist sense, whose results in many cases are reminiscent of the style of the great masters of this current in the history of painting.

Indeed, the Google Inceptionist algorithms, particularly DeepDream, share an intriguing parallel with Impressionist and Post-Impressionist art in their approach to visual representation. Both techniques emphasize the transformation of perception through their respective mediums – DeepDream by manipulating digital imagery to enhance patterns and features not immediately apparent, and Impressionism/Post-Impressionism by capturing moments of light and colour that elude the naked eye. DeepDream, akin to the way Impressionists break down scenes into individual colour strokes, deconstructs images into elemental patterns and textures, reassembling them into something new and, often, surreal. This process, much like the artistic movement, challenges traditional perceptions of reality, inviting viewers to reconsider the familiar through a distorted, dream-like lens that accentuates the underlying, often overlooked, aspects of visual stimuli.

If then one acts on the more superficial layers of the network, those deputed to the recognition of whole objects, then the result consists in a kind of artificial pareidolia:[5] if a neural network trained to recognize faces is presented with a photograph of clouds, the network will transform the image of the clouds so that human viewers too can recognize a face therein. Such a step is certainly interesting, and the images produced intriguing, but this is where the slippage between the doctor of neural networks and the augur of artificial intelligence takes place. Indeed, at a certain point, one gets the impression that

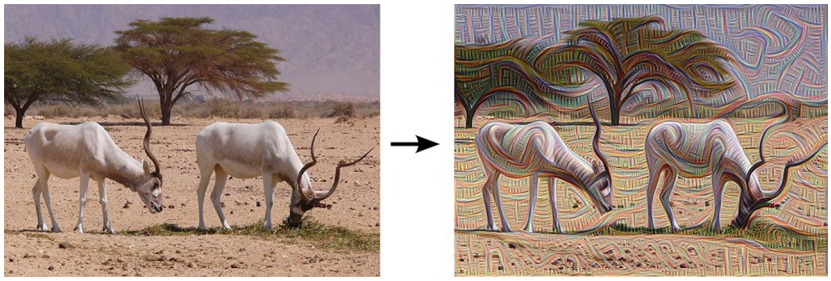

Figure 1. Original photo by Zachi Evenor. Right: processed by Günther Noack, Software Engineer. From the Google 'Inceptionism Gallery': https://photos.google.com/share/AF1QipPX0SCl7OzWilt9LnuQliattX4OUCj_8EP65_cTVnBmS1jnYgsGQAieQUc1VQWdgQ?pli=1&key=aVBxWjhwSzg2RjJWLWRuVFBBZEN1d205bUdEMnhB (accessed 29 May 2024; Licensed under CC BY 4.0: https://creativecommons.org/licenses/by/4.0/.

the authors of the post lose the precise sense of what they have been aiming at – which was trying to understand the functioning of the different layers of an artificial neural network trained to recognize the subject of a certain image – and essentially begin to fiddle with the network. The idea that you can discover interesting things by fiddling with the digital, and in particular with artificial neural networks, is quite popular. In fact, the ideological assumptions of this idea should be better explored.

It would perhaps be foolhardy to claim that Newton, Galileo, all the way to Einstein and contemporary science, have fiddled with their devices and computation until they arrived, as if by serendipity, at a useful result. The attitude of science, and especially that of modern science, has always been, on the contrary, to favour a prior consideration for method. One does not discover things by fiddling, but by refining the method by which one investigates the studied reality. Semiotics, which attempts to emulate the scientific disciplines, seeks to do the same. The attitude manifested in the Google post, on the other hand, is much more characteristic of disciplines such as architecture or engineering, which, though based on scientific thinking, value the heuristic nature of doing. The problem, however, is that here this doing – this fiddling with neural networks by putting them upside down, so to speak – is first presented as an heuristic method of meta-epistemological nature, suitable for better understanding the inner workings of the networks, while later it is in fact adopted to produce knockout images, the heuristic value of which, however, is unclear.

The techniques presented here help us understand and visualize how neural networks are able to carry out difficult classification tasks, improve network architecture and check what the network has learned during training. It also makes us wonder whether neural networks could become a tool for artists – a new way to remix visual concepts – or perhaps even shed a little light on the roots of the creative process in general.[6]

In fact, these images are no longer a technique but a rhetoric. They are like the augur's robe. They are put there to convey a sense of novelty, surprise, and even eeriness or uncanniness, but they serve no purpose, or at least they do not serve the effort of understanding that was promised at the beginning of the post. In it, in fact, many different discursive genres eventually mingle and fade into each other, contaminating each other; the authors present themselves as investigators of the heuristics of artificial neural networks, yet then go on to fiddle with their own experiment and, in the end, essentially close the blog as artists: reverse-engineering artificial neural networks is interesting because it produces intriguing images.

This third segment of the article has presented a compelling case study centred around the influential blog post 'Inceptionism: Going Deeper into Neural Networks' by Alexander Mordvintsev, Mike Tyka and Christopher Olah from Google. This exploration into the mechanics of artificial neural networks (ANNs) serves as an illustration of the interplay between technical proficiency and rhetorical presentation in the realm of digital technology.

The term 'Inceptionism', inspired by Christopher Nolan's film *Inception*, along with the metaphorical language used throughout the blog post (e.g. 'going deeper into neural networks'), underlines the importance of rhetorical strategies in the communication of complex technological concepts. Such metaphors not only make the subject matter more accessible to a broader audience but also imbue the technology with a sense of mystery and depth that exceeds its biological counterpart, thereby enhancing its perceived authority and significance.

The post details how ANNs are trained to recognize images by exposing them to millions of examples, adjusting their parameters to yield increasingly accurate results. However, the authors express a critical view towards the notion of 'understanding' these results solely based on their outward efficacy, pointing out the limitations of reverse-engineering as a means to truly grasp the inner workings of these networks.

One fascinating aspect of the case study is the technique of asking ANNs to transform an image full of visual noise into one that it recognizes, revealing cognitive patterns and biases developed during the learning phase. This approach, while initially aimed at demystifying the operational logic of ANNs, inadvertently transitions into a form of visual rhetoric that transforms ordinary images into artworks reminiscent of Impressionist or Post-Impressionist styles.

This shift from a scientific exploration to a form of artistic experimentation is emblematic of a broader trend where the boundaries between technique and rhetoric blur. Such endeavours, while producing visually intriguing results, may ultimately detract from the original goal of enhancing our understanding of ANNs. Instead, this trend underscores the increasing role of digital images as both a medium of scientific inquiry and a rhetorical tool designed to captivate and persuade.

FROM ARTIFICIAL HALLUCINATION TO HUMAN SPECULATION

Two questions emerge at this point. The first: what is really intriguing, surprising and fascinating about these images? The second: since they do not in fact say anything about the inner workings of the artificial neural networks that had produced them, then who will attribute purpose and meaning to them? It is here that the second, more radical shift from technique to rhetoric emerges: the humanities, and in this case especially cognitive psychology and art theory (Spratt, 2017), volunteer to illustrate the revelation of these images, even resorting to an archaeology of the mescalinic image (Klüver, 1926).[7]

It is necessary, first of all, to start with the tautological character of these visual experiments. The picture of clouds contains eyes because it has been retouched by an artificial neural network trained to recognize eyes. The most this visual experiment can reveal is how an image X must be transformed for it to appear to contain representations of an object Y. So far, this is really a triviality. The process through which this triviality can become the starting point for bold speculations in both cognitive science and art theory is complex and falls into an area necessarily subject to interdisciplinary study. First, an aspect related to the sociology of knowledge must be considered; in the digital age, universities and their public and private research centres continue to produce novelty and knowledge, but it is unquestionable that, in most cases, the most striking innovations in this field, and especially in that of computer vision and pattern recognition, come from the private laboratories of the world's major players in the production of technology in this field, namely Google, Apple, Meta, Envidia, etc. This fusion of the places and people who produce technology with the places and people who reflect on technology has important consequences that are known to anyone who has ever worked from the inside of one of these laboratories: the research that is conducted there is often free, curiosity-driven and blue-sky frontier, but it is still closely linked to a form of doing that is embedded in processes of technology production for the market. Consider how different the field of literary theory would be if its development had been delegated to publishers, or art theory if it had been developed by art dealers, or even theoretical physics itself if it had been developed by nuclear engineers. The *forma mentis* of those who research but always have in mind the possible practical application of their research is different from the *forma mentis* of those who do research purely for the sake of doing it, and in a context that explicitly and implicitly does not require that the gained knowledge has practical spin-offs (the borders can be porous, but there are borders).

The other element of the sociology of knowledge to consider is that these laboratories and their discoveries are only the starting point of a parasitic chain, so to speak, in which the added value created in them is capitalized and exploited to create added value in other fields, often quite remote both physically and disciplinarily from the source of these innovations. Underlying this parasitic chain is an element of sensationalism that is not entirely created

a posteriori but is part of the inherent logic by which these private laboratories produce and disseminate knowledge. The typical researcher in these centres wants to amaze. That is why they favour a certain kind of question, a certain kind of investigation, and a certain kind of answer. As for questions, one could say that those are favoured that awaken certain archetypes, but also inevitably certain stereotypes – some of them ancestral – in the relationship between human and nature through knowledge. Artificial neural networks, one is told, have risen to a rank of complexity that makes them almost creatures, which, though assembled by human creators, escape their control and thus become a mystery;[8] this mystery is also configured in a manner analogous to that which anthropologically grips humans, namely, the correspondence between exteriority and interiority, between what shows itself at the surface of the body and what instead takes place in its depths.

There is perhaps no adjective that has been more widely used in the field of artificial intelligence in recent years than 'deep'; and, indeed, it better than any other evokes an underwater topology of knowledge, in which there is a known surface but there are also abysses into which to dive in order to discover the secret, to grasp the truth, to gain the sense of how these neural networks really function inside. As will be seen, this myth of the mysterious depth to be plumbed, of remote anthropological origin, is in turn projected, or perhaps one should say 'introjected', toward knowledge of the depth from whose mystery it has in fact drawn its inspiration, namely that of the human brain and its workings; it will then be claimed, as some cognitive science scholars do, that understanding the depths of artificial neural networks helps one to understand those of biological neural networks.

This last segment has focused on the transition from the technical aspects of artificial neural networks (ANNs) to the rhetorical interpretations and implications of their outputs. This transition is marked by two pivotal questions: the intrinsic appeal of the images generated by ANNs and the attribution of meaning to these images when they do not necessarily reveal the internal mechanics of the ANNs themselves. The humanities, particularly cognitive psychology and art theory, step in to ascribe purpose and interpret these images, sometimes invoking comparisons to the effects of mescalin on visual perception as a way to understand the revelatory nature of these images (Clausberg 2010).

This analysis has revealed that the experiments producing such images are essentially tautological: an ANN trained to recognize eyes in images will, predictably, manipulate a cloudscape to reveal eyes. While this might seem trivial at first glance, it serves as a springboard for more profound interdisciplinary speculation in cognitive science and art theory about the nature and implications of these transformations. This speculation is rooted in a broader context where significant technological innovations, especially in computer vision and pattern recognition, often originate from the private research labs of major technology companies like Google, Apple and Meta. The research

environment in these labs, driven by curiosity and the pursuit of novelty, starkly contrasts with traditional academic research, highlighting a different mindset focused on practical applications and market-oriented outcomes.

This scenario underscores a 'parasitic chain' where the innovations from these labs are leveraged and repurposed across various fields, creating a cycle of knowledge and value extraction that extends far beyond the labs' physical and disciplinary boundaries. The inherent sensationalism in the dissemination of knowledge from these labs is not incidental but a deliberate aspect of their operation, aiming to astonish and captivate the audience. This approach to knowledge production and dissemination taps into deep-seated human archetypes and stereotypes about the natural world and our place within it, portraying ANNs as almost sentient entities that elude their creators' full understanding.

The widespread use of the term 'deep' in the context of artificial intelligence reflects a metaphorical evocation of knowledge as an ocean with unexplored depths, suggesting that penetrating the mysteries of ANNs can offer insights into the human brain's complexities. This metaphorical framing hints at a profound connection between the exploration of artificial and biological neural networks, proposing that understanding the former can illuminate the latter.

The transition from artificial hallucination to human speculation exemplifies the central hypothesis concerning the technical and rhetorical functioning of algorithms. It illustrates how the technical achievements of ANNs become subjects for rhetorical exploration and speculation, bridging the gap between the creation of algorithmic images and the human quest for understanding and meaning. This transition from technique to rhetoric and the subsequent speculation in the humanities highlight the complex interplay between technology and interpretation, underscoring the role of algorithms not only as tools for problem-solving but also as catalysts for intellectual and artistic inquiry.

A SLEIGHT OF HAND

As in a sleight of hand, however, and as is often the case in the rhetoric of mystery, there is something that escapes, a place and some actions that, while existing, are concealed with skillful illusionistic procedure. In the above post, for example, the authors admit that these neural networks contain 'black boxes', another metaphorical expression winking at a dimension of mystical secrecy as well as at the deep secret of human consciousness. They do not, however, reveal what they have done to create this mystery. The material or mathematical factory of these neural networks is not explained, nor is one described in detail how these networks were trained, with what images, and through what procedures. In a caricature comparison, one could say that the rhetoric of these kinds of articles is as follows: a new coffeemaker was invented

and it is used to make coffee but the coffee that comes out of it is really a mystery; one cannot understand how such a good beverage comes out. So, let one try to study in depth the characteristics of this coffee to understand how the coffeemaker makes it. Such curiosity is commendable, but would it not be more fruitful to try to satisfy it by beginning with the design of the coffeemaker and the nature of the blend that was used to make the coffee? In this, the technological production and market-oriented environment produces an obvious conditioning: the secret that is being attempted to be unveiled here is in fact not the ontological secret of scientists, or at least not totally thus. It is the industrial secret of the market. The rhetoric of mystery and revelation that, perhaps implicitly and unconsciously, the world's major technological production laboratories adopt, especially in the field of artificial intelligence, is as follows: we create artificial intelligence, but we do not tell you exactly how because this constitutes the added value of our work and the commodity we have to sell; yet we propose to you the idea that the results of this intelligence are inherently mysterious, and therefore deserve to be studied through indirect procedures that plumb the depths of the AI.

This critical reasoning does not exclude the fact that in the workings of AI there are indeed phenomena that are still poorly understood, emerging from the complexity of these systems and producing surprising effects; however, this reasoning points to the fact that studying these emergent phenomena without being able to know in depth the reality from which they emerge, and even shifting the focus to effects rather than causes, is not immune from a paradoxical industrial rhetoric, that of 'reverse-engineering'. Usually, this procedure is necessary to bypass an industrial secret, for example, as when the Chinese managed to create an aircraft carrier catapult by studying that of an Australian aircraft carrier they had purchased. But the aforementioned blog post proposes a paradoxical self-back-engineering, in which AI creators stage efforts to perform a back-engineering of what they themselves have created and whose trade secret they, or at least their company, guard.

A fundamental element enables the accomplishment of this sleight of hand. As in many rhetorical illusions, what makes it possible to distract and shift attention while creating the mechanism of revelation is a tool whose semiotic and rhetorical power is rooted in human neurophysiology and anthropology, and has been used since the dawn of time: images. This is not the occasion to dwell on the anthropological history of this mode of sign production and its ancestral appeal (Descola, 2021). Yet it is necessary to reflect on how images are used in research of this kind, both in its elaboration and dissemination. Indeed, the surprise effect generated by the dissemination of these investigations owes much to the presentation of digital images, which then enter a parasitic media ecology in which they are endlessly reproduced, re-released and, above all, decontextualized, that is, used for different purposes, detached from the cognitive intent with which their production was associated. Images attract the eye, they attract the mind, they attract the heart

and they also attract the hand to the wallet, albeit the virtual wallet of the digital marketplace. The specific case needs to be analysed in detail but it, too, must be placed in the context of a more general socio-semiotics of the contemporary iconosphere.

The metaphorical 'black box' of neural networks not only alludes to the mysteries of AI but also to the deep secrets of human consciousness, suggesting a layer of complexity and secrecy that goes beyond the technical into the mystical. The authors of the discussed post acknowledge these 'black boxes' but stop short of unveiling the processes that led to their creation, leaving the reader with a sense of wonder and perhaps frustration at the lack of transparency.

This rhetorical strategy, as highlighted in the segment, is akin to introducing a novel coffeemaker without explaining its inner workings, focusing instead on the quality of the coffee it produces. This approach underscores a market-driven motive where the intrigue surrounding the product's operation enhances its perceived value. A more fruitful pursuit of understanding would involve examining the design and material composition of the 'coffeemaker' (in this case, the AI systems), rather than just marvelling at the output. The real mystery being guarded, indeed, is not a scientific or ontological one but an industrial secret, tightly held by corporations to maintain a competitive edge in the market.

As regards the paradoxical practice of 'reverse-engineering' within the AI field, where creators of AI systems engage in attempts to backtrack and understand their creations, a process usually aimed at uncovering industrial secrets, it is presented as a strategic move to retain control over the proprietary knowledge while engaging in a form of intellectual exploration.

Central to this 'sleight of hand' is the role of images, which are powerful semiotic and rhetorical tools that captivate attention and stimulate curiosity. These images, produced or highlighted by AI research, enter a media ecosystem where they are endlessly circulated, often stripped of their original context and purpose, serving various ends that range from academic inquiry to commercial exploitation.

That exemplifies the intricate balance between the technical and rhetorical aspects of AI research. It illustrates how the rhetoric of mystery and the strategic withholding of information contribute to the aura surrounding AI technologies, influencing both public perception and academic discourse. This balance between revealing and concealing, between the allure of the unknown and the pursuit of knowledge, mirrors the broader theme of the present article: the relationship between the tangible advancements in AI (the technical dimension) and the narratives constructed around these advancements (the rhetorical dimension). This dynamic interplay shapes our understanding of AI, prompting a deeper reflection on the implications of such technologies beyond their immediate functionalities, into the realms of philosophy, ethics and socio-economic impact.

CONCLUSIONS

In contemporary society, and especially in the digital sphere, there is rarely any value creation that is not accompanied and indeed even sanctioned by image creation. The development of image-making technology and its widespread dissemination has certainly fostered this trend, but it has been coupled with a much longer tendency: iconophilia, but also iconocentrism, in which this mode of sign production, with all that distinguishes it in terms of semiotic, phenomenological and anthropological dynamics, is posited as central to the enunciation of meaning and its transmission among human beings. Iconocentrism essentially says: 'look, there is an image, therefore there is meaning'. Of course, there is always meaning in an image, but this meaning is not necessarily always the most relevant and pertinent in the semiotic exchange between two or more individuals (D'Armenio, 2022). The rhetorical power of the image is immense, yet its semiotic effect does not always lead to revealing the essentials of a communicative interaction.

Take as an example academic communication itself, for instance, that of lectures in a Philosophy Department such as the one in which the present author works. Until much of the 1990s at least, no philosophy professor would have dreamed of showing pictures during lectures; at most, in some areas, such as aesthetics, 35 mm slides were used. These were common especially among art historians, who bought them from the museums of Europe or had them purchased through their university libraries or photo libraries. A good art history department in Europe, still throughout the 1990s, had to be equipped with slide projectors. Between 21 and 27 July 2002, I attended a congress of IAWIS, the International Association for Word and Image Studies, in Hamburg, and I remember that the participants, among whom were many art theorists and art historians, were still delighted to be able to use the same slide projector that had been used by Erwin Panofsky. Indeed, for a university lecturer in much of the 1990s who wanted to add images to their verbal discourse, the possibilities were essentially two: slides – which were very popular precisely in the field of art theory – and transparencies – which, on the other hand, were used mainly by economists. I still remember the transparencies that my Professor of Political Economy at the University of Siena, Maurizio Franzini, would pre-package and then complete while projecting them by means of the special device. I also remember that one day he forgot to place the transparency on the projector and drew a graph on the same glass surface of the lamp, moreover with indelible markers. Even Umberto Eco, when talking to his students about comic books or other images of popular culture in the 1980s and much of the 1990s, did not reproduce them but described them. Economic disciplines, earlier than others, actually introduced transparencies as a visualization aid to academic discourse, where the content displayed was mainly diagrams, that is, graphs. This practice was related to the world of business, where the presentation of diagrammatic images, and especially graphs,

to support verbal discourse had become the practice, especially by means of whiteboards or panels. University lecturers until much of the 1990s, on the other hand, were immersed in an essentially logocentric semiotic *Umwelt*, where they could also talk about film, television, contemporary art or whatever but they did not show it, rather they described it and, in most cases, merely evoked it.

In fact, the first PowerPoint presentation from a computer-connected laptop took place in Paris on 25 February 1992 by Robert Gaskins, the program's principal creator. Here is how Gaskins himself describes the event:

> The very first public use of a laptop to project video from PowerPoint took place on February 25, 1992, at the Hotel Regina, in the Place des Pyramides, Paris (across from the Tuileries). With a laptop casually under my arm, I entered at the back of a ballroom filled with hundreds of Microsoft people from the European, Middle Eastern, and African subsidiaries. I walked through the audience carrying the laptop, up to a podium at the front; there I opened the laptop, and plugged in a video cable on the lectern. I began delivering a presentation to introduce PowerPoint 3.0 for Windows, using PowerPoint 3.0 running on the laptop feeding video out to a projector the size of a refrigerator which put the 'video slides' onto a huge screen behind me. No one had ever seen PowerPoint running on a portable computer before, let alone being used to produce a real-time video show in color with animated builds and transitions. The audience, all Microsoft people who talked to customers frequently, grasped immediately what the future would bring for their own presentations; there was deafening applause. (https://www.robert-gaskins.com, accessed 29 May 2024).

Robert Gaskins had foreseen in the preparatory plans that, thanks to the spread of personal computers and with the gradual miniaturization of projectors, these presentations would move from large specialized business meeting rooms to small rooms for business meetings. What could not yet be foreseen is instead ably described by Gaskins in a following passage of the same biographical note:

> All this was predicted in my strategy documents from the mid-1980s; what was unexpected was that the same hardware would also extend PowerPoint use into university teaching, children's school reports and science fair projects, sermons in churches, super-titles for opera houses, and many other uses that its creators had never imagined.

Today, most of my colleagues in my Philosophy Department would never lecture without the aid of a PowerPoint. At the last International Semiotics Congress, held in Thessaloniki between 30 August and 4 September 2022, I did not listen to any paper that did not make use of PowerPoint. There has been

a shift from a phase in which the use of this program was useful for adding value to presentations through the possibility of showing pictures to a phase in which, in order for a presentation not to be devalued and belittled, it must be accompanied by PowerPoint. Teaching in China without the aid of PowerPoint, for example, is now considered a lack of professionalism. In Italian telematic universities, underpaid young researchers are essentially hired to produce PowerPoint in bursts for their students. Only in the rarefied world of snobbish academic high culture can one afford not to present with images, and that is a privilege reserved for the few. Yet one does not have to be a sophisticated visual semiotician to realize that, in many cases, images are not the subject of communication, as was the case with art historians with their slides – which they required in black and white anyway precisely so that the colour would not distract from the attribution exercise – and neither do they add relevant content to what is being talked about. Images, on the other hand, are used for what could be called pragmatic purposes, that is, to exercise a phatic function that is not only one of contact with the interlocutor/observer, but also with the community of interpreters in which one communicates, with the common sense in which one circulates content, a community and a common sense that now regard images as indispensable for maintaining a nexus of attention in the passage of information between sender and receiver.

The rhetorical slippage that occurs when the attempt to 'back-engineer' a neural network that recognizes images leads to the production of fabricated images through artificial intelligence, is also accomplished by situating itself in this long period of the emergence of images, and especially digitally produced images, as the place where the pragmatic eroticism of contemporary communication is produced, that which incites the recipient of a communication to want to see more, and therefore hear more, and receive more content. This slippage, however, also causes the added value of research produced in the major AI laboratories on the planet to be revived and parasitized primarily in the domain of images.

In this article, we have explored the interplay between algorithmic processes and their cultural implications, revealing the multifaceted nature of algorithms as both technical constructs and rhetorical entities. By analyzing the semiotics of algorithmic images, we uncovered the nuanced ways in which these images serve as sites of negotiation between human and computational intelligences, reflecting and shaping societal norms and values. Through the lens of 'inceptionism', we have probed the paradoxes inherent in the visualization of algorithmic processes, highlighting the tensions between transparency and opacity in our understanding of algorithmic functioning. This article underscores the importance of a critical perspective towards the epistemological and ontological questions raised by the increasing pervasiveness of algorithms in our digital culture. It calls for a reevaluation of the ways in which visual rhetoric and algorithmic authority co-construct our perceptions of reality, knowledge, and power in the digital age. By critically engaging with the implications of

algorithmic authority, we pave the way for a more nuanced understanding and interrogation of the digital landscapes that shape our contemporary existence.

ACKNOWLEDGEMENTS

This article results from a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No 819649–FACETS); an initial version of the article was presented in the form of a talk entitled 'The Augmented Face: Digital Mask or Digital Veil?', within the symposium 'The Augmented Image', University of Liège, 6–7 October 2022. I thank the two organizers of the symposium, Maria Giulia Dondero and Enzo D'Armenio, as well as those colleagues who provided feedback after the presentation. I am also in debt to the two anonymous reviewers of the first submitted version of the article. The final version of the article was completed thanks to time and work conditions provided by FRIAS, the Institute for Advanced Studies at Freiburg University, Germany, thanks to the European Union's Horizon 2020 research and innovation programme under the Marie Sktodowska-Curie grant agreement no. 754340.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

ORCID ID

Massimo Leone https://orcid.org/0000-0002-8144-4337

NOTES

1. The strategy of augurs does not solely rely on deceit; it is possible for an augur to genuinely believe in their practices, meaning that they are not deliberately trying to mislead others; that does not change the fact that their strategies of prediction have no fundament whatsoever, although they might firmly believe in them.
2. It is not necessary to possess expertise in divination or analogous rituals to acknowledge that such practices were not conducted in an arbitrary vacuum. Rather, they are believed to have been integral to a comprehensive worldview or a distinct body of knowledge. This perspective does not serve to advocate for divination but to recognize that these practices were inherently connected to a particular cosmology, possibly of an analogical nature, following the framework proposed by Philippe Descola. This connection might elucidate the persistence of such practices over time, despite their deviation from

the principles of contemporary technique and science, suggesting a nuanced interpretation is warranted.

3. Minimalist rhetoric refers to a strategic approach in discourse where authority is constructed by deliberately reducing or subtracting signs, rather than by adding more. This method emerges as a response to criticism that exposes the superficial or pseudo-technical nature of traditional, more elaborate rhetorical strategies. Minimalist rhetorics rely on simplicity and the omission of excess to convey authenticity and credibility, countering the scepticism fostered by critical analysis. An example of minimalist rhetoric can be found in modern advertising, where brands may use straightforward, no-frills messaging to differentiate themselves from competitors who employ more elaborate and embellished claims, aiming to present themselves as more genuine, trustworthy, or focused on essential values.
4. Available at: https://blog.research.google/2015/06/inceptionism-going-deeper-into-neural.html?_gl=1*1urce7z*_ga*MTYzNzA1MzY5OC4xNzA3NzEyMjEw*_ga_163LFDWS1G*MTcwNzcxMjIwOS4xLjAuMTcwNzcxMjIxNS4wLjAuMA (accessed 29 May 2024).
5. Pareidolia is a psychological phenomenon where people perceive recognizable patterns or objects, such as faces or animals, in unrelated and random stimuli. This can happen with visual cues, like seeing shapes in clouds or hearing hidden messages in music when played in reverse. It is a form of apophenia, which is the tendency to attribute meaning to perceived connections or patterns between unrelated things.
6. Available at: https://blog.research.google/2015/06/inceptionism-going-deeper-into-neural.html?_gl=1*1urce7z*_ga*MTYzNzA1MzY5OC4xNzA3NzEyMjEw*_ga_163LFDWS1G*MTcwNzcxMjIwOS4xLjAuMTcwNzcxMjIxNS4wLjAuMA (accessed 29 May 2024).
7. As an example of this 'stepping in', consider Emily L Spratt's paper, 'Dream formulations and deep neural networks', which delves into the alignment of deep learning technologies with art historical and cognitive psychology theories, particularly through Google's *DeepDream* and Georgia Tech's *Grad-CAM*. Spratt argues for the significant insights these technologies provide into image recognition processes, drawing parallels to human perception and learning. She advocates for the potential benefits of integrating these machine learning tools into art historical research, suggesting they can help decode the 'black box' of image recognition and enhance our understanding of visual categorization. However, this perspective invites critiques regarding the potential for oversimplification of human cognitive processes, the risks of reinforcing biases through technology and the implications of heavily relying on proprietary technologies in academic research.

Spratt's enthusiasm for cross-disciplinary collaboration underscores an optimistic yet contentious view of the intersection between technology and the humanities, highlighting both the potential and the pitfalls of incorporating AI into the study of art history.

8. 'Mystery' refers to a phenomenon or subject that elicits intrigue and curiosity due to its unknown or incomprehensible nature. In the context described, it emerges from the complex interactions and outcomes of artificial neural networks that, despite being creations of human intellect, surpass their creators' understanding and control. This elusiveness fosters a sense of wonder and enigma, drawing parallels to the anthropological fascination with the unseen forces or processes that underpin observable realities, encapsulating both the allure and the ambiguous essence of the unknown.

REFERENCES

Annus A (2010) *Divination and Interpretation of Signs in the Ancient World*. Chicago, IL: Oriental Institute of the University of Chicago.

Bertrand, Denis. 2023. "'Le visage vu à l'envers est méconnaissable'", 21–34. In Giuliana, Gianmarco Thierry and Leone, Massimo, eds. 2023. *Sémiotique du visage futur* ("I saggi di Lexia", 50). Rome: Aracne.

Cary P (2008) Before words were signs: Semiotics in Greek philosophy, online. In: Cary P (ed.) *Outward Signs: The Powerlessness of External Things in Augustine's Thought*. Oxford: Oxford Academic, online edn. DOI: 10.1093/acprof:oso/9780195336498.003.0003 (accessed 17 May 2023).

Clausberg K (2010) Warburg und Meskalin: Argumente aus dem Schweigen, online (1–28). *Kunstgeschichte: Open Peer Reviewed Journal*. Available at: https://www.kunstgeschichte-ejournal.net/180/ (accessed 29 May 2024).

D'Armenio E (2022) The rhetorical dimension of images: Identity building and management on social networks. *Semiotica* 246: 87–115. DOI: https://doi.org/10.1515/sem-2020-0113.

Descola P (2021) *Les formes du visible*. Paris: Éditions du Seuil.

Dondero MG and Fontanille J (2012) *Des images à problèmes: Le sens du visuel à l'épreuve de l'image scientifique*. Limoges: Presses universitaires de Limoges.

Greimas AJ (1984) Sémiotique figurative et sémiotique plastique, online. *Actes sémiotiques* 60. Available at https://www.unilim.fr/actes-semiotiques/3848 (accessed 17 May 2023).

Klüver H (1926) Mescal visions and eidetic vision', 502–515. *American Journal of Psychology* October 37(4): 502–5015. DOI: 10.2307/1414910.

Leone, Massimo. 2011. "Rituals and Routines: A Semiotic Inquiry", 107–120. *Chinese Semiotic Studies*, 5, 1; DOI: https://doi.org/10.1515/css-2011-0109

Leone, Massimo. 2013. "The Semiotics of Fundamentalist Authoriality", 227–39. In Leone, Massimo, ed. 2013. *The Authoriality of Religious Law, special issue of the International Journal for the Semiotics of Law*, 26, 1; ISSN: 0952-8059; DOI https://dx.doi.org/10.1007/s11196-012-9304-7

Leone, Massimo. 2015. "Double Debunking: Modern Divination and the End of Semiotics", 433–77. *Chinese Semiotic Studies*, 11, 4. DOI: https://doi.org/10.1515/css-2015-0025

Leone, Massimo. 2021. "Indexes: Cultural Nature and Natural Culture", 95–112. *Rivista di estetica*, 76, 1; ISSN: 2421–5864; DOI: https://doi.org/10.4000/estetica.7699

Leone, Massimo, ed. 2013. "Preface", 91–94. In Leone, Massimo, ed. 2013. The Authoriality of Religious Law, special issue of the *International Journal for the Semiotics of Law*, 26, 1; DOI https://dx.doi.org/10.1007/s11196-012-9298-1

Leone, Massimo and Santangelo, Antonio, eds. 2023. *Semiotica e intelligenza artificiale* ["Saggi di Lexia", 48]. Roma: Aracne.

Lotman J (2019) Technological progress as a culturological problem. In: Tamm M (ed.) *Juri Lotman: Culture, Memory and History*. Cham, Switzerland: Palgrave Macmillan, 201–223. DOI: 10.1007/978-3-030-14710-5_15; originally published as Tekhnicheskii progress kak kul'turologicheskaia problema. *Trudy po znakovym sistemam* 22: 97–116, trans. Tamm M from Lotman I (2000) *Semiosfera*. St Petersburg: Iskusstvo–SPb, 622–638.

Manetti G (1987) *Le teorie del segno nell'antichità classica*. Milan: Bompiani.

Michaelstadter C (1913) *La persuasione e la rettorica*. Genoa: Formiggini.

Mordvintsev A, Tyka M and Olah C (2015) Inceptionism: Going deeper into neural networks. Available at: https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html (last accessed 17 May 2023).

Nguyen A et al. (2016) Plug & play generative networks: Conditional iterative generation of images in latent space, online (1–35). *ArXiv:1612.00005v2*. DOI: 10.48550/arXiv.1612.00005.

Nguyen A, Yosinski J and Clune J (2016) Multifaceted feature visualization: Uncovering the different types of features learned by each neuron in deep neural networks, online (1–23). *ArXiv:1602.03616v2*. DOI: 10.48550/arXiv.1602.03616.

Spratt EL (2017) Dream formulations and deep neural networks: Humanistic themes in the iconology of the machine-learned image, online (1–29). *ArXiv:1802.01274*. DOI: 10.48550/arXiv.1802.01274.

Vaan M de (2008) *Etymological Dictionary of Latin and the Other Italic Languages* (Leiden Indo-European Etymological Dictionary Series 7). Leiden: Brill.

Wei Donglai et al. (2015) Understanding intra-class knowledge inside CNN, online (1–7). *ArXiv:1507.02379*. DOI: 10.48550/arXiv.1507.02379.

BIOGRAPHICAL NOTE

MASSIMO LEONE is Professor of Philosophy of Communication, Cultural Semiotics and Visual Semiotics at the Department of Philosophy and Educational Sciences, University of Turin, Italy, part-time Professor of Semiotics in the Department of Chinese Language and Literature, University of Shanghai, China, associate member of Cambridge Digital Humanities, University of Cambridge, Director of the Institute for Religious Studies at the 'Bruno Kessler Foundation', Trento, and Adjunct Professor at the University of Caracas, Venezuela He has been Visiting Professor at several universities in the five continents. He has single-authored 16 books, edited more than 60 collective volumes and published more than 600 articles in semiotics, religious studies and visual studies. He is the winner of a 2018 ERC Consolidator Grant and of a 2022 ERC Proof of Concept Grant. He is editor-in-chief of *Lexia*, the *Semiotic Journal of the Center for Interdisciplinary Research on Communication*, University of Turin, Italy, co-editor-in-chief of *Semiotica* (De Gruyter), and co-editor of the book series 'I Saggi di Lexia' (Aracne), 'Semiotics of Religion' (Walter de Gruyter) and 'Advances in Face Studies' (Routledge). Since 2024, he is a member of the Academy of Europe in the section of Art History and Musicology.
Address: University of Turin, Via Sant'Anselmo 8, Turin 10124, Italy. [email: massimo.leone@unito.it]