
aracne

Classificazione Decimale Dewey:

005.8 (23.) SICUREZZA DEI DATI

GERARDO **IOVANE**, GIOVANNI **IOVANE**, DARIO **SILVESTRI**

ARTIFICIAL INTELLIGENCE AND HYPNOTIC COMMUNICATION

aracne

ISBN
979-12-218-2522-0

PRIMA EDIZIONE
ROMA 18 FERBBRAIO 2026

Summary

FOREWORD

Why this book was written

This volume is the result of the intersection of three converging trajectories: the evolution of generative artificial intelligence systems, the established tradition of hypnotic and Ericksonian communication, and the growing need to understand how these fields intertwine in ways that are as promising as they are delicate. In recent years, the advent of large language models (LLMs) has radically transformed the landscape of technology-assisted communication. These are no longer simple text processing tools, but systems capable of generating linguistic, visual and audio content with a fluidity and semantic coherence that are strikingly reminiscent of some of the dynamics typical of hypnotic communication [1].

Suggestion, understood as the ability of language to direct attention, modulate states of consciousness and facilitate experiential changes, has always been one of the most powerful tools in human communication. Milton H. Erickson demonstrated how the very structure of language through metaphors, ambiguities, pacing and leading can induce perceptual and cognitive changes without necessarily resorting to formal protocols of hypnotic induction [2]. Today, generative AI systems are capable of replicating many of these linguistic patterns with unprecedented accuracy and scalability. This raises fundamental questions: how does the nature of suggestion change when it is mediated by an algorithm? What opportunities open up for professionals in communication, coaching and psychotherapy? And, above all, what ethical responsibilities arise when such powerful tools become accessible on a large scale?

The underlying hypothesis supporting this work is that artificial intelligence should not be understood as a digital hypnotist, but rather as a cognitive amplifier: a system that enhances human communication skills, making the creation of suggestive content more efficient, while keeping creative, ethical and clinical control in the hands of the human professional. Generative AI, in fact, lacks intentionality, genuine empathy, and the ability to adapt to relationships in real time—all essential qualities in any intervention involving

suggestion and the modification of states of consciousness. However, its ability to rapidly generate text, images, and audio content consistent with specific communicative intentions makes it an extraordinary tool for designing immersive experiences, hypnotic scripts, guided visualisations, and psychoeducational materials.

This book, therefore, stems from the need to provide a conceptual and operational map for those who wish to explore this new territory with scientific rigour and ethical awareness. It is not a technical manual for AI developers, nor is it a clinical text reserved for psychotherapists: rather, it is an interdisciplinary work that engages with psychology, cognitive science, computer science, and communication science, offering theoretical and practical tools for understanding and using generative AI in contexts of hypnotic and suggestive communication. The ideal audience includes coaches, trainers, professional communicators, psychologists and counsellors interested in integrating these technologies into their work, as well as researchers and students who wish to understand the implications of this emerging convergence between technology and the mind.

Personal and professional motivations

Our professional experience lies at the intersection of clinical practice, training in hypnotic communication and the exploration of emerging technologies. For years, we have been working with language, suggestion and states of consciousness, both in therapeutic and training contexts, and we have observed the evolution of artificial intelligence systems with growing interest. When, in 2022, models such as GPT-3.5 and later GPT-4 began to demonstrate unprecedented linguistic capabilities, it became clear that a new phase was opening up in the relationship between technology and persuasive communication. These were no longer simple chatbots or virtual assistants: these systems were capable of producing complex narratives, articulate metaphors and fluid semantic progressions, all of which are central features of hypnotic communication.

The main motivation for writing this book was the realisation of a gap: while the literature on generative AI was growing rapidly in the technical and commercial spheres, there was a lack of rigorous and in-depth analysis of its implications in the field of suggestion and hypnotic communication. There was, and continues to be, a wealth of popular content, often commercial in nature, promising miraculous applications of AI in psychology and

education, but without a solid theoretical basis or adequate ethical reflection. At the same time, the academic and clinical world seemed uninterested or even wary of these technologies, sometimes dismissing them as mere automation tools with no scientific value.

This polarisation seemed insufficient and unproductive to us. Generative AI, if understood and used competently, can represent a significant extension of human communication skills, accelerating creative processes, personalising content on a large scale and making hypnotic communication techniques accessible to professionals who would otherwise not have the resources to develop them independently. At the same time, it is undeniable that these tools carry real risks: the possibility of unconscious manipulation, the creation of addiction to immersive content, the spread of pseudoscientific practices masquerading as technological innovation [3]. Writing this book was, therefore, a way to address these tensions explicitly and constructively.

On a personal level, engaging with generative AI has also been a stimulating intellectual challenge. Exploring the capabilities of these systems has forced us to spell out many of the implicit assumptions that guide our communication practice. Questions such as what makes a metaphor effective, how a hypnotic progression is structured, and which linguistic elements promote attentive absorption have become central not only to communicating with AI, but also to refining our theoretical understanding of suggestion. In this sense, AI has also functioned as an epistemological mirror: a technology that has allowed us to reflect more consciously on the mechanisms we use daily in our work.

Why AI radically changes the way we use suggestion

Generative artificial intelligence introduces a paradigm shift in the production and dissemination of suggestive content. Traditionally, the creation of hypnotic scripts, guided visualisations, therapeutic metaphors or immersive narratives required advanced linguistic skills, narrative creativity and a deep understanding of the psychological mechanisms underlying suggestion. These processes were inherently artisanal: each professional developed a personal repertoire of techniques, images, and communication patterns over time, often through years of practice and supervision. Generative AI, on the other hand, enables the rapid, large-scale production of complex linguistic content,

drastically reducing creation times and allowing for a level of personalisation that was previously unthinkable.

One of the most relevant elements is the ability of generative models to produce seamless semantic continuity and narrative coherence. An LLM can develop an extended metaphor while maintaining a constant semantic field, modulate the rhythm of a text according to precise instructions, and systematically integrate elements of linguistic pacing (mirroring the client's language). This fluidity represents a significant advantage over manual writing, which requires continuous cognitive effort to maintain internal consistency and hypnotic progression. In technical terms, generative models operate through probabilistic prediction mechanisms based on enormous bodies of text, which makes them particularly effective at reproducing recurring linguistic patterns, including those typical of hypnotic communication [4].

A second crucial aspect concerns multimodality. While in the past the creation of hypnotic content was mainly limited to written text or recorded voice, today AI allows for the simultaneous generation of text, images, audio and even video. This opens up unprecedented possibilities for the construction of immersive experiences that integrate multiple sensory channels, amplifying the overall evocative effect. A concrete example: it is possible to design a guided visualisation session in which the hypnotic text is generated by AI, converted into audio with a modulated synthetic voice to promote states of relaxation, and accompanied by symbolic images created using text-to-image models. All this can be achieved in a few hours, where in the past it would have required weeks of coordinated work between writers, professional speakers and designers.

However, the most profound change is not technological, but conceptual. Generative AI forces us to rethink the role of the professional: no longer (just) a direct creator of content, but a designer of communication processes, a curator of experiences and an ethical guarantor of quality and safety. The fundamental shift is from execution to strategy: while AI takes care of material production, the professional defines objectives, chooses narrative structures, calibrates suggestive intensities, and verifies clinical and contextual consistency. This reconfiguration of work requires new skills: the ability to formulate effective prompts (prompt engineering), to critically evaluate the outputs generated, to

integrate human and automatic contributions harmoniously, and to always maintain a reflective stance with regard to the limits and risks of automation [5].

Finally, AI democratises access to tools that were once reserved for the few. A coach with basic training in hypnotic communication can now create professional-quality materials without necessarily possessing advanced skills in creative writing or audio production. This raises important ethical questions: on the one hand, it increases the spread of potentially beneficial practices; on the other, it lowers barriers to entry, making it possible for sensitive techniques to be used improperly or superficially. It is precisely for this reason that we believe a rigorous and informed training approach is essential.

Ethical and legal warnings

This book is not a clinical manual and in no way replaces specialist training in psychotherapy, clinical hypnosis or counselling. The techniques and tools described are intended for professionals working in the fields of communication, coaching and personal growth, and must always be used within the limits of their professional competence and in compliance with current regulations. The use of suggestion, even when mediated by technology, implies a direct responsibility towards the people with whom one works. The ability to influence states of consciousness, direct attention and facilitate experiential change requires not only technical competence, but also ethical integrity, awareness of one's own limitations and a willingness to constantly supervise one's own work.

On a legal level, it is essential to note that in many countries the use of hypnosis in a therapeutic setting is regulated and reserved for licensed healthcare professionals (doctors, psychologists, psychotherapists). Even when hypnotic techniques are not formally used, but one works with suggestive language, guided visualisations or immersive content, it is necessary to always operate within one's professional competence. A coach cannot practise psychotherapy; a trainer cannot treat clinical symptoms; a communicator cannot replace a healthcare professional. AI, however powerful, does not change these boundaries: on the contrary, it makes them even more relevant, as it facilitates the creation of content that could exceed the limits of competence if used without discernment [6].

On an ethical level, key issues related to informed consent and transparency arise. When using AI-generated content in coaching, training or psychological support contexts, it is important that those involved are aware of the nature of these tools. It is not a question of hiding the use of AI, but of contextualising it appropriately: clients have the right to know that the materials used have been created with the support of automated systems, even though they have been selected, personalised and validated by a human professional. This principle of transparency is consistent with international guidelines on the ethical use of AI in psychology and training [7].

Another ethical dimension concerns cognitive vulnerability. People who turn to professionals for personal growth, habit change or emotional support often find themselves in states of particular openness and receptivity. The use of AI-enhanced suggestive techniques amplifies the communicative power of the professional, but also their responsibility. It is essential to avoid any form of manipulation, to ensure that content is always geared towards the person's well-being, and to maintain an attitude of deep respect for the client's autonomy and dignity. Ethical suggestion is that which facilitates awareness, not dependence; which promotes autonomy, not conditioning; which respects the person's timing and limits, without coercion.

Finally, it is important to emphasise that generative AI has structural limitations that must always be considered. Current models lack deep semantic understanding, genuine empathy, real-time relational adaptability, and sensitivity to non-verbal cues that are fundamental to any authentic human interaction. AI can produce technically correct but emotionally inappropriate texts, formally coherent but clinically inappropriat r metaphors, and seemingly empathetic but relationally inauthentic content. For this reason, human review is always mandatory, and professional judgement must remain central to every stage of the process. This book is written on the assumption that AI is a tool, not a substitute; an amplifier, not a decision-maker; a support, not a professional. And like any powerful tool, it requires competence, responsibility and constant ethical vigilance.

CHAPTER 1: Introduction

1.1 What is meant by hypnotic artificial intelligence

The term hypnotic artificial intelligence requires a precise definition to avoid conceptual and applicative misunderstandings. We are not referring here to AI systems capable of inducing hypnotic states in the clinical sense, nor to machines with autonomous hypnotic capabilities. Rather, we intend to describe the conscious and strategic use of generative artificial intelligence systems for the creation, optimisation and distribution of communicative content that incorporates principles, structures and patterns typical of hypnotic and Ericksonian communication. In other words, AI becomes a tool at the service of human professionals to amplify the suggestive effectiveness of communication, while maintaining design, ethical and relational control in the hands of the operator.

Hypnotic communication, in the Ericksonian tradition, is not limited to the formal induction of trance, but includes a wide repertoire of linguistic techniques aimed at facilitating states of focused attention, imaginative absorption and openness to change [8]. These states are not necessarily altered in a strong sense, but rather represent particular modes of information processing, characterised by reduced critical activity, greater receptivity to indirect suggestions and easier access to imaginative and symbolic resources. Hypnotic language operates through mechanisms such as pacing (mirroring subjective experience), leading (guiding towards new experiences), the use of transformative metaphors, productive ambiguity, semantic redundancy and controlled narrative progression [9].

Generative artificial intelligence models, particularly Large Language Models, have structural characteristics that make them particularly suitable for producing this type of language. Firstly, their architecture based on attention mechanisms allows them to maintain semantic consistency over long narrative arcs, generating texts that present thematic continuity, logical progression and stylistic fluidity. This is essential for the creation of hypnotic scripts, which require gradual development, controlled redundancy and the absence of narrative discontinuities that could interrupt attentive absorption [4]. Secondly, LLMs are trained on huge bodies of text that include, among other things,

literary narratives, poetic texts, meditative content and psychoeducational materials: all areas in which the evocative dimension of language is central.

A concrete example can clarify the concept. Suppose we want to create a guided visualisation script for a stress management coaching programme. Traditionally, the professional would have to write the text manually, paying attention to elements such as sentence rhythm, the use of sensory imagery, progression from the concrete to the abstract, the integration of permissive suggestions (you may notice..., perhaps you notice...) and the construction of a narrative that gradually leads to a state of relaxation. Using a generative AI system, it is possible to provide structured instructions (prompts) that guide the model in generating text with these characteristics. The prompt could specify: a calm and reassuring tone, the use of natural metaphors (water, trees, mountains), progression from visual images to bodily sensations, the inclusion of suggested pauses, and permissive rather than directive language. The AI then produces a first draft, which the professional reviews, personalises, and adapts to the specific context of the client.

It is essential to emphasise that AI does not perform hypnosis: hypnosis is a relational process that requires presence, real-time adaptation, reading of non-verbal signals and empathic attunement: all dimensions that current systems do not possess. What AI can do is generate communicative materials that incorporate hypnotic principles, which will then be used by the professional in appropriate relational contexts. Hypnotic artificial intelligence is therefore a hybrid practice: it combines the computational ability to generate coherent and evocative language with the human expertise to design, contextualise, personalise and ethically validate this content. It is a tool for enhancement, not replacement.

1.2 Why the term is metaphorical and not clinical

The use of the term hypnotic in this context is deliberately metaphorical and requires important clarifications to avoid pseudoscientific drifts or inappropriate applications. In scientific and clinical literature, hypnosis is defined as an interactive process in which a qualified professional uses specific procedures to facilitate states of consciousness

characterised by focused attention, reduced peripheral awareness and increased responsiveness to suggestions [10]. This process is inherently relational: hypnotic effectiveness depends on the quality of the therapeutic relationship, the clinician's ability to adapt in real time to the client's responses, and the integration of verbal and non-verbal communication. An AI system, however sophisticated, cannot replicate this intersubjective dimension.

When we talk about hypnotic artificial intelligence, we are referring instead to the application of linguistic and communicative principles derived from the hypnotic tradition to the design of AI-generated content. It is a practice that falls within the realm of persuasive communication, coaching, psychoeducation, and the creation of immersive experiences, but it does not replace clinical hypnosis nor can it be equated with it. The metaphor is useful because it highlights the structural similarities between certain linguistic patterns produced by AI and those used in hypnotic communication (semantic continuity, redundancy, metaphors, narrative progression), but it would be misleading to interpret it as a statement of functional equivalence.

One of the main differences concerns the dimension of dynamic adaptation. In an authentic hypnotic interaction, the practitioner continuously observes the client's responses: micro-movements, changes in breathing rhythm, changes in facial expression, variations in tone of voice – and modulates their communication accordingly. This real-time feedback loop is essential for calibrating the suggestive intensity, modifying metaphors if they do not resonate, and slowing down or speeding up the pace according to the client's state. AI-generated content, however well designed, is static: it cannot adapt to individual responses, cannot perceive signs of discomfort or overload, and cannot modify its communication strategy on the fly [11].

A second critical aspect concerns intentionality and ethical responsibility. Clinical hypnosis is always oriented towards the client's well-being, set within a defined therapeutic context, guided by agreed objectives and subject to professional supervision. AI, on the other hand, has no intentionality: it generates content based on statistical patterns learned during training, without any understanding of the meaning or implications of what it produces. This means that an apparently hypnotic text generated by AI could

contain inappropriate suggestions, inadequate metaphors or narrative progressions that, in certain contexts, could be counterproductive or even harmful. This is why human review is always mandatory: only a competent professional can assess the clinical, ethical and contextual appropriateness of suggestive content.

Clarifying that the term is metaphorical also serves to prevent commercial misunderstandings. There is a real risk that the expression 'hypnotic AI' will be used in a sensationalist way to promote dubious applications, such as apps that promise to automatically hypnotise users or services that boast therapeutic capabilities without any scientific basis. This book takes a critical and rigorous stance precisely to distinguish itself from these abuses: AI can be a valuable tool for competent professionals, but it is not a magic shortcut and cannot replace training, experience and human responsibility. The term hypnotic is an interpretative lens, not a promise of clinical performance.

1.3 Communication, attention and imagination

At the heart of the encounter between artificial intelligence and hypnotic communication are three fundamental cognitive processes: communication, attention and imagination. Understanding how these processes interact is essential for using AI effectively and responsibly. Communication, in the broad sense, is not simply the transmission of information, but the co-construction of shared meanings through symbolic, linguistic and relational processes. When we communicate, we do not transfer content from one mind to another, but create interpretative contexts that allow the interlocutor to construct mental representations consistent with our communicative intentions [12]. In hypnotic communication, this process is intentionally oriented towards facilitating specific experiential states: relaxation, concentration, imaginative immersion, openness to change.

Attention is the selective mechanism that determines which stimuli are processed as a priority and which are filtered out. Neuroscientific research has shown that attention is not an unlimited resource: we have limited attentional capacity, and our mind operates through continuous selection processes [13]. Hypnotic communication works strategically with these mechanisms: it narrows the field of attention to specific stimuli (a voice, a

mental image, a bodily sensation), reducing peripheral awareness and creating favourable conditions for absorption. Generative AI systems can contribute to this process by creating linguistic content that facilitates attentional focus: controlled rhythm, semantic redundancy, gradual progression, use of recurring linguistic anchors.

A practical example: suppose we want to create a guided meditation audio. The AI-generated text can be structured in such a way as to progressively narrow the focus of attention. It starts with broad and inclusive suggestions (notice the sounds around you...), moves on to more specific dimensions (bring your attention to your breath...), and guides you towards internal experiences (imagine a quiet place...). This progression is not random, but reflects established principles of attention management: from the external to the internal, from the concrete to the abstract, from the sensory to the imaginative. AI, when properly trained, can replicate this structure in a consistent and systematic way.

Finally, imagination is the ability to construct mental representations of scenarios, experiences and situations not directly present in immediate perceptual reality. In the Ericksonian tradition, guided imagination is one of the central tools: through narratives, metaphors and detailed sensory descriptions, it facilitates access to experiential states that can have therapeutic, formative or transformative effects [14]. Generative AI excels at producing imaginative narratives: it can create elaborate stories, describe multisensory scenarios, and develop extended metaphors while maintaining symbolic coherence. This is possible thanks to the ability of LLMs to draw on vast narrative repertoires and combine them in original but semantically coherent ways.

The synergy between these three processes—communication, attention, imagination—is what makes AI-assisted hypnotic communication effective. Communication provides the vehicle (language, images, sounds); attention determines the degree of absorption and focus; imagination creates the experiential space in which change occurs. When these elements are orchestrated consciously and competently, AI becomes a powerful amplifier: it accelerates production times, increases stylistic quality and allows for sophisticated customisation. However, without a deep understanding of these cognitive processes, there is a risk of producing content that is technically correct but

psychologically ineffective, or worse, content that overloads the attentional system, creates imaginative confusion, or induces states of discomfort rather than well-being.

1.4 Interdisciplinarity of the text

This book is built on a deliberately interdisciplinary conceptual framework. The convergence between artificial intelligence and hypnotic communication cannot be adequately understood by remaining within a single disciplinary field: it requires the integration of perspectives from psychology, cognitive science, computer science, neuroscience, communication science, and applied ethics. This methodological choice responds to an epistemological necessity: complex phenomena require multiple conceptual tools, and the use of AI in sensitive communicative contexts such as hypnotic ones is certainly a complex phenomenon.

From a psychological point of view, we draw mainly on the Ericksonian and post-Ericksonian tradition, which has provided sophisticated descriptive models of the mechanisms through which language influences subjective experience. Authors such as Milton Erickson, Ernest Rossi, Stephen Gilligan, and Jeffrey Zeig have documented in detail how specific linguistic patterns (presuppositions, implications, double binds, isomorphic metaphors) facilitate cognitive and emotional changes [15]. This tradition is particularly relevant because it does not merely describe hypnosis as an altered state, but conceptualises it as a communicative process: a perspective that aligns perfectly with the use of AI as a tool for communicative design.

Cognitive science offers us theoretical models that are essential for understanding attentional processes, working memory, semantic processing, and mental imagery. The r literature on selective attention, cognitive load, mental representations, and embodied simulation provides an empirical basis for understanding why certain linguistic content is more effective than others in facilitating specific experiential states [16]. These contributions are fundamental in avoiding intuitive or folkloric approaches: what works in hypnotic communication is not magic, but reflects well-documented cognitive mechanisms that can be intentionally activated through appropriate communicative choices.