Proceedings from the Document Academy

Volume 11
Issue 2 *Proceedings from the 2024 Annual Meeting of the Document Academy*

Article 10

2024

Documents in the Age of Non-Human Agency

Bernt Ivar Olsen-Kristiansen Dr
Microsoft, berntivarolsen@outlook.com

Niels W. Lund
UiT Arctic University of Norway, niels.windfeld.lund@uit.no

Recommended Citation
Olsen-Kristiansen, Bernt Ivar Dr and Lund, Niels W. (2024) "Documents in the Age of Non-Human Agency," *Proceedings from the Document Academy*: Vol. 11 : Iss. 2 , Article 10.
DOI: https://doi.org/10.35492/docam/11/2/10
Available at: https://ideaexchange.uakron.edu/docam/vol11/iss2/10

From the very beginning of human life, documentation is an essential part of daily life as part of a complementary process of documentation, communication and information (Lund 2024). To communicate with the environment, to be informed by the environment, humans need to make or experience a document, like a scream by a baby, a vocal discrete entity coming up as fast as it is dissolved. Documents are made and experienced all the time in many different ways by using many different means.

Often documents are understood as primarily a physical object, but it is important to stress that any document is not only a physical object, but just as much a social and mental object. If we take Napoleon's hat we may consider it as a document about the emperor (Pagès 2021 and Day, 2024), to be shown in a museum, but if we didn't know that it had been used by Napoleon, it has just been a hat and not been a unique document and may have been a document in textile history or costume history. When archeologists do excavations finding black postholes in the ground, they may claim the postholes to be documents of former houses. In both cases, remains, a hat and post holes, are used as means in the documentation process together with interpretational texts in articles, books, exhibitions etc. Through history, humans have used and developed machines, instruments and tools to support human documentation.

In the last 40–50 years, the computer and digital technology in general has played a growing role in all kinds of documentation processes in all parts of society. The so-called Artificial Intelligence (Simon, 1996) has in the last 10–15 years been pushing the borders between what is considered human-made and machine-made. Nevertheless, it is important to stress what Simon says, that "Artificial things are synthesized by human beings" (Simon, 1996, 5) In other words, artificial intelligence, (AI), are man-made and not made by the machines themselves.

In this paper, we will explore the newest development in AI, the LLMs and Diffusion models and discuss how the products of these systems can be understood as documents.

Have we reached the point where machines can express themselves autonomously, or are their output merely statistical or stochastic reproductions of human expressions? We will examine the technical nature and features of these models, their general capabilities, and provide contextual examples of what is achievable with this technology. Finally, we must also consider the social role and implications of these LLM programs.

Let us first briefly establish the current understanding of "machine agency" and explore the foundational concepts of Large Language Models that underpin recent advancements such as ChatGPT, Claude, Perplexity, and others. Additionally, we will discuss "Diffusion models," which are the basis for AI Image generators (e.g., text-to-image) like Dall-e, Stable Diffusion and others.

Generative AI

Generative AI, often referred to as Gen AI, is a branch of artificial intelligence focused on creating original content such as text, images, video, audio, or software code in response to a user's prompt or request. This technology relies on sophisticated machine learning models, particularly deep learning models, which simulate the learning and decision-making processes of the human brain. These models identify and encode patterns and relationships in vast amounts of data, enabling them to understand natural language requests and generate relevant new content.

Generative AI has a wide range of applications across various domains, including digital art, content creation, software development, and more. By leveraging these advanced techniques, generative AI can produce novel and creative outputs that mimic human creativity and intelligence.

Let's dive into two notables of these that enable generation of human-like content, LLMs and diffusion models.

Large Language Models (LLMs)

Large Language Models (LLMs) are advanced AI systems designed to understand and generate human language. Fundamentally, they operate by predicting the next token (or word) in a sequence, utilizing architecture such as Transformers. Transformers are advanced AI models that process and understand sequential data by using an "attention mechanism", allowing them to comprehend entire contexts rather than analyzing information piece by piece. Unlike traditional models that translate or process language word by word, Transformers can simultaneously examine whole passages, understanding complex relationships between different elements. They work like intelligent language detectives, rapidly learning patterns and capturing nuanced meanings across various domains such as translation, text generation, and conversational AI. By giving different weights to relevant parts of the input, Transformers can generate remarkably human-like and contextually accurate outputs, revolutionizing how machines understand and interact with language. Moreover, these models are trained on vast and diverse datasets, enabling them to capture subtle language patterns and contextual meanings. During training, LLMs repeatedly predict the next token, learning to anticipate the following word based on the context of preceding words (He, Hangfeng et al. 2024). This process involves multiple layers of computation, where each layer refines the model's understanding of the relationships within the text, thereby improving accuracy over time.

LLMs convert the input (text data) into "embeddings," which are vector representations of words or tokens. These embeddings pass through a series of layers, with each layer enhancing the model's comprehension, identifying both immediate and broader contexts. This layered processing allows LLMs to surpass simple surface-level predictions, capturing syntax, semantics, and deeper linguistic structures.

Emerging Capabilities

As LLMs continue to expand in scale, they exhibit capabilities far beyond mere next-token prediction. Larger models, trained on increasingly comprehensive datasets, demonstrate the ability to generate coherent text, recognize context, translate text, and even perform tasks requiring reasoning. This means that

LLMs can generate entire paragraphs of coherent thought, making them useful for tasks such as summarization, translation, and content generation.

A key innovation is the use of *multi-token prediction* (He et al., 2024), which allows LLMs to predict several tokens simultaneously rather than one at a time. This improves both the efficiency and accuracy of the models, making them more practical for real-world applications. This advancement enables faster and more reliable processing of large bodies of text, such as interviews, survey responses, or open-ended questions, thus allowing humans to concentrate on interpretation and theory development rather than manual data processing.

As noted, LLMs are trained on vast datasets tend to capture deeper principles of language, often aligning with linguistic theories. For instance, LLMs can exhibit structural patterns akin to *power-law distributions* (He et al., 2024), which are observed in natural language usage. These patterns reveal underlying social or cultural phenomena, providing new perspectives on how language reflects societal norms, values, or inequalities. Some prominent examples of LLMs in products we have become familiar with are ChatGPT, Claude AI, Gemini and Perplexity as well as the open source model Llama.

Diffusion Models

Diffusion models are a type of AI that generates data by gradually transforming random noise into structured, meaningful content (Chen et al., 2023; Wikipedia, 2023). They are especially powerful for creating visual art and images. Let's break down how they work and what they can do.

Initially, a piece of structured data like an image is gradually turned into random noise by adding layers of Gaussian noise. Think of this as slowly blurring a clear picture until it becomes unrecognizable. The model then works backward, step-by-step, to remove the noise and recreate the original image. A neural network, trained specifically for this task, helps clean up the noise bit by bit. This allows the model to generate new, clear images from pure noise.

New Modes of Expression Enabled by Diffusion Models

Diffusion models can create images that look like hand-drawn sketches. By controlling the amount of noise removed at each step, the model can produce various styles, from rough sketches to detailed drawings. These models can generate stunning digital paintings and artwork. Whether it's mimicking famous artists' styles or creating entirely new artistic expressions, diffusion models offer a versatile tool for digital artists. They can also create images that look incredibly real, almost like photographs. This capability is useful for generating high-quality images for advertising, design, and visual media without needing a camera. Beyond still images, diffusion models can generate sequences of images that form short movie clips. By carefully controlling the transformations over time, the models can create smooth transitions between frames, making it possible to produce animated sequences or video content.

One of the most exciting applications is generating images from text descriptions. For example, given a description like "a sunset over a mountain range," the model can produce a visually accurate and compelling image based on the textual input. Like text-to-image, diffusion models can generate video clips from text descriptions. This involves creating a sequence of coherent images that flow seamlessly to form a short video. For instance, describing a "bird flying across the sky" can result in a short, animated clip of that scene. These models are excellent at filling in missing parts of an image. If parts of an image are corrupted or missing, diffusion models can predict and fill in the gaps, making the image whole again. Diffusion models can enhance the resolution of low-quality images, adding detail and clarity. This is particularly useful in applications requiring high-definition visuals, such as medical imaging or satellite imagery. Given noisy or corrupted data, diffusion models can clean up the noise and restore the original quality. This is useful for improving the clarity of images, audio files, or any other type of data affected by noise. These models can apply specific artistic styles to images, transforming a normal photo into a piece of art with the style of a famous painter, for example.

Artists can use diffusion models to explore new creative avenues, generating unique artworks and experimenting with different styles. These models can

produce visual effects, animated sequences, and realistic images for movies, games, and virtual reality. High-quality, realistic imagery can be generated for product promotions, saving time and resources in photoshoots.

Diffusion models have revolutionized the way we create visual content, offering new modes of expression that range from hand-drawings and paintings to photorealistic images and movie clips. By transforming random noise into detailed and coherent visuals, these models open endless possibilities for artists, designers, and content creators. They are typically used in applications such as text-to-image, text-to-video, image inpainting, super-resolution, denoising, and style transfer, making them versatile tools in various creative and professional fields.

The conceptual, technical- and computational/hardware developments and inventions that have given us different kinds of Generative AI such as LLMs and diffusion models and their likes as outlined above together provide us with the basis of our discourse for what implications this might have for the role of these new tools and agents in the creation of documents. Let us dive deeper into this in the following sections.

Intention and agency

In our previous work we have outlined our view of a document model as a complementary model that encompasses both processes of documentation and communication as well as information (Lund, 2024). We have also outlined an ontology of human expression (Olsen et al., 2012) based on this model that encompasses (at least) six constituents of a document:

- The agent(s)
- The means
- The modes
- The cognitive components
- The social components
- The physical components

In the current age of Artificial Intelligence, the advent of Large Language Models (LLMs) marks a significant shift in how documents and documentation are conceived and created. As discussed, documents have historically served as a medium for human expression, bridging communication across time and space (Lund, 2024). The insertion of computers and information systems into this process has undoubtedly altered the means and modes of documentation, yet the fundamental concept of a document has remained unchanged (Olsen et al., 2012; Lund, 2024).

The integration of Gen AI into our documentation model challenges the notion that documents are solely human expressions. While these models operate based on statistical patterns derived from vast human-created texts, their ability to generate novel and contextually appropriate content suggests a new dimension of document creation. This leads us to explore whether AI can be considered as entities capable of self-expression or if their output remains fundamentally tied to human input and interpretation.

Intentionality vs. Agency in Document Theory

Both intentionality and agency are central to how we understand human activities, particularly in the creation of documents. Let's begin by examining the relationship between the two and how agency might evolve in this context. Intentionality in the traditional sense refers to the purposeful act of creating a document to communicate, inform, to teach or to 'evidence' something. In human terms, intentionality is intrinsically tied to conscious thought, decisions, and goals. However, AI-generated outputs complicate this model, as the outputs themselves lack the conscious, goal-oriented nature of human activities. What we see instead is a probabilistic generation of text that mimics intentionality, but without the underlying mental state driving it.

Agency typically refers to the capacity of an entity to act autonomously, make choices, and affect its environment. In human terms, this involves decision-making, responsibility, and accountability. But how do we think about agency in the context of AI, especially in relation to document creation?

Floridi and Sanders (2004) introduce the concept of Levels of Abstraction (LOA) to analyze artificial agency, arguing that an entity's status as an agent

depends on the specific LOA at which it is considered. They suggest three criteria for agenthood at a given LOA: interactivity (the ability to respond to stimuli), autonomy (the ability to change state without stimuli), and adaptability (the ability to change transition rules based on experience). This framework resonates with what we will denote *distributed agency* in the context of AI and document creation. While we acknowledge that AI systems, particularly LLMs, lack full autonomy and conscious intentionality, their ability to generate novel and contextually appropriate outputs suggests a form of functional agency

Human–AI Collaboration and Distributed Agency

When we introduce AI into the documentation process, the agency involved in creating a document shift from a purely human-centric model to one including shared or distributed agency.

Distributed Agency*:* In the context of document creation, LLMs operate as part of a larger network of actions that includes human input (the prompt), the design of the AI system (created by developers), and the AI's processing capabilities. Here, the "agent" is no longer a single entity but the collaboration between human and machine, similar to one between human agents, albeit with the exceptions mentioned above, such as e.g. missing accountability from the non-human agent.

One possible framework for understanding AI's role in documentation one could view the human-AI interaction as one of shared or hybrid agency. Instead of treating LLMs as passive tools, we can think of them as semi-autonomous collaborators within a network of human agents. Each element of the system contributes to a form of agency, though the ultimate accountability lies with the human participants. This would require redefining the nature of agency in terms of influence, where AI tools influence the document creation process without assuming full control or accountability.

When viewed through Floridi and Sanders' lens, AI systems, at an appropriate LOA, exhibit interactivity through their responses to prompts, autonomy through their complex internal processing, and adaptability through machine

learning – although the latter may involve re-training of the LLM or fine-tuning it. AI systems, though not fully autonomous agents, can participate in document creation as collaborators within a network of distributed agency, contributing to the evolving nature of documentality in the digital age. They act within the constraints of their training data and the algorithms governing them, as well as the structural constraints like the server park that host them and us as humans instructing them. Unlike human agents, who can reflect on their actions and make deliberate choices, LLMs are typically reactive systems. They do not choose to act; they generate responses based on probabilistic weights and prompts (the latter being human intentions).

Agency in humans also involves accountability—humans are responsible for their actions and the consequences that follow. AI, particularly LLMs, cannot be held accountable in this way. Any responsibility for AI outputs falls back on the human agents involved in their development, deployment, and usage. AI, therefore, lacks moral agency – unless it is instructed to adhere to ethical rules and guidelines. In the latter case that would be a human agency, nonetheless.

Corporate and Organizational Agency in the Context of LLMs

In addition to individual human agents and machine agents, corporate and organizational agencies add another layer of complexity to the interaction with Large Language Models (LLMs). These institutional agents, such as companies like Microsoft and OpenAI, act as both mediators and shapers of human-AI interaction. Unlike individual human agents or LLMs, corporate agents operate with collective intentionality, making structured decisions that fundamentally shape the development, deployment, and governance of AI systems. This agency, however, is distinct from both individual and machine agency due to its institutional nature, which involves a network of actors functioning under shared objectives.

Corporate agents wield significant influence through their ability to shape the LLM's architecture and capabilities. This agency manifests various critical aspects, including design choices that determine the behavior and limitations of LLMs, the selection and curation of training data, the implementation of safety measures and ethical guidelines, and the control over access and

deployment. Corporate agents also establish policies and restrictions that directly impact how these systems are used and by whom. Although individual employees contribute to these decisions, the corporate entity itself exercises a distributed form of unified agency that shapes the broader AI ecosystem.

Regulatory and Governmental Agency

Another important form of organizational agency in the LLM landscape is that of governmental and regulatory bodies. These institutions exercise agencies by creating and enforcing policies, establishing legal frameworks that govern AI development and use, setting safety and ethical standards, and overseeing market competition. Their influence operates at a meta-level, not only determining how LLMs function but also shaping the corporate agents' capacity to exercise their agency in developing and deploying these systems.

This interplay between corporate and governmental agencies forms a complex web of influences, constraints, and enablement that ultimately defines how individual humans interact with LLMs. While corporations make choices about how LLMs are built and used, government bodies ensure that these decisions align with broader societal goals, such as fairness, safety, and innovation. Together, corporate and governmental entities create a structured framework that governs both the technical and ethical dimensions of AI, situating LLMs within a broader socio-technical system of accountability and agency.

The artificial nature of LLMs and their agency

To give further context to the capabilities of current AI, of which there are quite substantial discussions in the various fields of Artificial Intelligence, we can refer to the recent Nobel laureate, Geoffrey Hinton in one of his recent lectures (University of Toronto – Geoffrey Hinton, October 2023)

Geoffrey Hinton stands in opposition to AI skeptics like Yann LeCun and Noam Chomsky, asserting that Large Language Models (LLMs) demonstrate genuine understanding of language. At the core of his argument is the belief that accurate next-word prediction requires real comprehension, much like

answering questions demands understanding their meaning. This isn't mere statistical pattern matching, but rather sophisticated processing that enables true understanding.

While LLMs do perform a type of "autocomplete," their function goes far beyond simple text matching. Instead, they learn through complex interactions between features using numerical activations, creating a nuanced understanding of language rather than just storing and retrieving text. This sophisticated processing enables them to demonstrate remarkable reasoning capabilities, including handling complex concepts like time and causality, solving logical problems, and even performing analogical reasoning—traditionally considered a uniquely human trait.

What critics often label as "hallucinations," Hinton reframes as "confabulation," drawing a parallel to human memory processes. Just as humans reconstruct memories rather than retrieve them perfectly, LLMs exhibit similar behavior. He points to real-world examples like John Dean's Watergate testimony, where human memory showed similar imperfect reconstruction patterns, suggesting this trait indicates sophisticated understanding rather than a flaw.

The foundation of LLM understanding, according to Hinton, lies in their ability to learn vast networks of feature interactions rather than storing static text. These interactions form complex rule systems that explain language patterns, with modern transformer architectures improving on early models by better handling ambiguity. This creates a dynamic form of understanding that contrasts sharply with traditional AI's static representations. Meaning emerges from the interplay of features, shaped by context in ways remarkably similar to human understanding.

Hinton’s view above is heavily contrasted with that e.g. of Noam Chomsky, who states that

> *“[…] machines don't do anything. I have a computer in front of me it's basically a paperweight doesn't do anything [..] all the computer in front of me is capable of doing is implementing a program. What's a program? Well a program is a theory written in*

> *a notation in which machines can implement […]"* (Chomsky, N. October 2023)

Chomsky critiques the idea of attributing agency or intentionality to machines, highlighting that machines are merely tools that execute human-designed programs. This perspective challenges the notion of AI autonomy, suggesting that AI systems are fundamentally dependent on human instructions and goals. Chomsky is in good company within a whole group of researchers within AI, such as Yann LeCun – a colleague of Hinton and a Turin award winner for his work within AI.

Adding to these perspectives we can consider the philosopher and linguist Baudrillard and simulacra (Baudrillard, 2020) to complement our discourse as Gen AI creates a world increasingly defined by simulations. These technologies, designed to imitate human communication, mirror his notion of hyperreality—where the boundary between the real and the representation collapses.

Imagine entering a theme park like Disneyland. Baudrillard argued that it exists not to immerse visitors in fantasy, but to convince us that life outside its gates is real by contrast. In truth, Disneyland's hyperreality blurs the line, rendering all of society a kind of theme park. Now consider LLMs as a linguistic version of this phenomenon. Just as Disneyland reinforces the idea that reality exists beyond the park, AI-generated texts—carefully marked as artificial—may bolster the illusion that human-generated content remains authentic. Yet both forms of expression, Baudrillard would argue, are simulations, where distinctions like real versus artificial lose significance. As noted by HA Simon in his famous treatise: artificial simply means human-made (Simon, H. A. 1996) (paraphrased).

Coladangelo (2022) explores how Disney theme parks, through meticulous design and immersive storytelling, construct a hyperreality that blurs the boundaries between the real and the simulated. Just as Disney invites visitors to suspend disbelief and engage with its carefully crafted illusions, AI systems, through their sophisticated algorithms and vast datasets, generate outputs that simulate human expression. This creates a new kind of hyperreality within

the human-machine interface, particularly in the co-authoring process. Human intentions, expressed through prompts, interact with the complex inner workings of the AI, resulting in documents that, while demonstrably artificial, can possess qualities of coherence, creativity, and even a simulated form of emotional intelligence.

This interplay between human input and AI processing leads to a hyperreal space where the lines between creator and creation, intention and outcome, become increasingly blurred. Coladangelo's analysis of Disney's constructed authenticity further highlights this dynamic. Disney theme parks meticulously curate a sense of authenticity by incorporating real-world references and historical motifs while simultaneously acknowledging their fabricated nature. Similarly, AI-generated documents often draw upon vast repositories of human knowledge and creative expression, creating a sense of borrowed authenticity while simultaneously revealing their artificial origins. This blurring of boundaries challenges traditional notions of documentality, prompting us to reconsider how meaning and intentionality are negotiated within these technologically mediated environments.

This idea gains further weight when we think about how LLMs generate not just language but understanding itself. These models work from patterns of patterns, creating outputs that, while syntactically correct, bear no direct relationship to an original "real." The result is a hyperintelligent system, producing language without reference to the human experience that typically grounds it.

As LLMs and other forms of generative AI continue to advance, they create hyperreal environments—texts that are often more coherent, articulate, and persuasive than human writing. In this way, they may not only reflect but generate hyperreality, blurring the lines between human and machine communication so thoroughly that, one day, distinguishing between the two may become irrelevant. LLMs don't just simulate language; they simulate understanding, knowledge, and perhaps even consciousness. They are trained on representations of representations (text and images about text and text and images about reality). Baudrillard might have considered the LLM as a realistic simulation of human expression. The LLMs are extremely good

simulators. They are one advanced step in the long history of simulations in the history of mankind

Machine agency in the document model

In "What Kind of Documents Can AI Language Model Outputs Be?" Donner (2024) explores the potential for AI-generated outputs, particularly from Large Language Models (LLMs), to be classified as documents within the field of documentation theory. Donner introduces and later updates the Model of Documentation Activity (MoDA), which is a framework for understanding the act of creating documents in physical, mental, and social contexts.

The revised model proposed by Donner (2024), named MoDA2, incorporates LLM outputs as a novel type of data provider, distinct from traditional human testimony and nature. A central argument is that AI outputs represent a new form of data—*artificially blended testimony*—created through LLMs' processing of vast corpora. These outputs are different from traditional testimony, as they lack direct human intention and are produced by algorithms rather than human experiences. The boundaries between human and machine roles, he argues (ibid.), in the documentation process are blurring, raising questions about authorship, intention, and agency in document creation.

In Donner's reflection on whether LLM outputs can be considered documents, a key issue is intentionality. Traditionally, a document requires intentional creation to convey information or evidence. LLM outputs, however, lack this kind of human intention—they are algorithmically produced based on probabilistic models trained on data. Donner hints that the human act of prompting AI could provide some indirect intentionality. We agree that the concept of intention may need special consideration. We would also like to consider the notion of task as part of the intention: "What is the task?"

One could argue that intentionality, as a concept, may need expansion in document theory. As we increasingly rely on autonomous systems, outputs that derive meaning and coherence from indirect human inputs (such as prompting) could be considered intentionally emergent. This would suggest

a new layer of intentionality where the prompt serves as a meta-intention guiding the AI's action without dictating specifics. On the other hand, one can also consider prompts as paratext (Skare, 2021): With Gen AI the paratext is all the metadata, the context, the way we prompt them, and even the user interface we interact with.

For those already regularly interacting with LLMs they will recognize how prompting for a *desired output*, whether this is a LLM producing some text or a image- or video-generator AI, this is a process of trial and error and - where proficiency and knowledge about the workings of the AI impacts the quality of the output; much similar to how a painter is dependent on their skill set and materials to output their desired output and expression. The prompts – the instructions used – as well as a notation of the specific models applied in the creative process can thus be considered paratext – or an intrinsic part of the document, if required.

Our exploration of intentionality in the context of AI-generated documents aligns with, but also expands upon, Robert Pagès' (2021) observations on the intentional nature of documentation. Pages emphasizes how documents are intentionally created and curated for interpretation, a process that involves human agency in selecting and shaping symbolic representations. We build upon this by examining how AI, particularly LLMs, introduces new layers of complexity to this process. While AI systems themselves lack the conscious intent of human creators, their outputs can be seen as "intentionally emergent," reflecting the meta-intentions embedded in human prompts. This concept resonates with Pagès' observation that even objects not originally intended as documents become so through deliberate acts of conservation and interpretation. However, we go further by proposing a model of distributed intentionality, acknowledging the combined influence of designers, data curators, and users in shaping the outputs of AI systems. This expanded view of intentionality recognizes the collaborative nature of documentation in the age of AI, where human intentions interact with the complex processes of machine learning to produce meaning.

Moreover, our discussion of distributed intentionality in the context of AI-generated documents also connects with the broader concept of

documentality, as explored by Buckland (2014). Buckland argues that documentality extends beyond traditional notions of documents to encompass any object "considered as" evidence. This could be an object purposely "made as" a document, "made into" one through contextualization, or even a natural object, like a leek, interpreted as symbolic evidence. Buckland emphasizes the role of context and interpretation in determining an object's documentary function, a perspective that aligns with Pagès' emphasis on the intentionality behind document creation. Pagès highlights the enduring symbolic representation embedded within a document, whether intentionally created or preserved for future interpretation.

Ronald Day (2024) further connects these ideas, arguing that Pagès' concept of the "auto-document," an object that inherently transmits information about itself, foreshadows contemporary discussions of documentality as a philosophy of information. For example, a fossilized bone becomes a document about a prehistoric creature through scientific interpretation, and Napoleon's hat serves as an "auto-document" by embodying historical significance. However, the emergence of AI-generated content, like text from LLMs or images from diffusion models, complicates these established frameworks. While traditional documentality relies on human agency and intentionality in shaping and interpreting symbolic forms, AI systems, trained on vast datasets, introduce a new layer of distributed intentionality, where human prompts interact with complex algorithmic processes to produce meaning.

In other related work, Tim Gorichanaz's (2015) argument for the co-created nature of documents aligns with our discussion of AI's role in document creation. Gorichanaz emphasizes that a document is not merely an information object but rather a product of the interaction between an information object and a human being. He argues that human perception, interpretation, and context are crucial in establishing document status, suggesting that documents reside in a "psycho-physical or socio-physical realm". This perspective sheds light on how AI, by entering the authorship process, fundamentally alters the dynamics of co-creation. AI systems, like LLMs, do not simply process information; they actively shape and contribute to the meaning-making process. As AI systems become more integrated into

the documentation process, the human element of perception and interpretation becomes intertwined with the AI's computational capabilities. This creates a new layer of co-creation where human intentionality, expressed through prompts, merges with the AI's ability to generate novel and contextually relevant content, resulting in documents that reflect the combined influence of both human and machine agency. This raises crucial questions about the evolving nature of authorship, intentionality, and the boundaries of documentality in the age of AI.

All the above reflections raise crucial questions about the nature of agency in document creation: can we ascribe agency to AI systems that lack conscious intent? Do they function as mere tools, or as collaborators in a hybrid form of agency? These questions demand a reevaluation of documentality in the age of AI, recognizing the blurring lines between human and machine in the creation and interpretation of documents, extending the idea of distributed agency from the human realm into a collaborative human-machine realm.

We build on the notion of distributed intentionality along with distributed agency. In general, within documentation one needs more focus on distributed agency/intentionality in all kinds of documentation processes, like the case of books with author, printer, publisher etc. having different tasks to do in the complex process of book publishing or any case of documentation. If one uses the notion of task, one may consider what kind of task is the LLM performing – for the human? What is it that the human agent wants the machine to help them express?

On a related note, there is another interesting intersection between human and machine as we are talking about documents as means to show, to instruct or to *teach*. In fact, one of the very prominent uses of LLMs today are as teachers – as instructors of knowledge that humanity has accrued up until now. One could argue that machine has become the instructor – and is now teaching humans – documenting as they go. On the other hand, one can argue that this teaching is based upon all the common knowledge that was previously only in the minds of experts, then (also-) in books and written documents as well and later into digital form, along with all documented forms of teaching. Before LLMs we assigned the document agency to the authors of

the books, or digital forms and perhaps those who 'instructed' with the aid of that knowledge, teachers and other tutors. Now we could assign that agency to the LLM that can now teach anything from a toddler to a scholar about pretty much anything known – in a personalized way, and in any language that the model has learned. However, even though LLMs may in this case have the agency of teaching it still needs to be prompted to teach – necessitating an intentionality to learn by a human in the first place.

Implications for Document Theory

As AI has progressed significantly over the past decade, to the point where we discuss their potential agency in documentation and whether – or at what point they might become conscious. The rapid acceleration over last couple of years, starting with the advent of ChatGPT, has had us question what consciousness is – where there in fact is no good current criteria or definition that can help us evaluating the state of current AI. As machines approach the existing measures of intellect to that of a human, we also start questioning what it means to be human – let alone what a document is – and what it means to document. Already accepting that agency can be distributed across humans, organizations and societies, we can begin to reconsider how documents are conceptualized considering Gen AI:

• Co-Creation of Documents: In the traditional model, documents are seen as the product of an individual or collective human agent's intentional act. If AI systems are viewed as possessing some form of instrumental agency, the creation of documents could be framed as a co-creation process with *distributed agency*. AI outputs, though generated by non-human actors, influence the final form of the document, suggesting that the document is the result of both human and machine agency working in tandem.
• Documents as Products of Agent Networks: Rather than viewing documents as the output of a single, intentional human act, we might start to see them as the result of an agentic network where human and machine actors collaborate. The document becomes a node in a larger network of interactions, with both human and AI inputs contributing to its creation.

Challenges and Future Considerations

While this expanded understanding of agency offers valuable insights into AI-human collaboration, it also raises several philosophical and ethical questions:

- Accountability in Co-Creation: If documents are co-created by humans and AI, how do we assign responsibility for errors, biases, or misleading information? While human agents ultimately control the AI and its usage, AI systems increasingly generate complex and unpredictable outputs that can blur the lines of accountability.
- Trust in Distributed Agency: As we move towards models of shared agency, how do we maintain trust in the output generated by AI? Humans are often inclined to trust outputs from LLMs because of their coherence, but we must remain critical of their limitations, particularly in terms of accuracy, bias, and transparency.
- Moral and Ethical Considerations: AI outputs are increasingly influencing decision-making in fields such as law, medicine, and journalism. As AI becomes more integrated into the documentation process, we must address the ethical implications of relying on systems that lack moral agency yet influence outcomes in morally significant ways.

As we have explored, LLMs challenge traditional conceptions of both agency and documentality. They operate not as fully autonomous agents, but as participants within a network of distributed agency. This reflects a significant departure from the historical notion of documents as purely human expressions, mediated by physical or digital means. In this context, LLMs complicate our understanding of intentionality, especially when we consider the evolving role of humans as prompt engineers and co-creators.

The intentionality of a document has traditionally been rooted in the conscious decisions of its creator. However, with LLMs, we enter a space where intentionality becomes distributed across the human-AI interface. The prompt supplied by a user serves as a meta-intention that directs the AI's output but does not determine its specific form or content. This raises critical questions: To what extent can we ascribe intentionality to the final output, and is it

possible to view this as an emergent property of human-machine interaction? Note also that the prompt is currently typical text or in some cases image(s) or even sound but increasingly it is also multi-modal, just like our human senses.

We must also address the concept of agency. While LLMs lack cognitive autonomy, they exhibit a form of "functional agency," contributing to the document creation process in ways that resemble human participation. This challenges us to reconsider the rigid boundaries between tool and agent, pushing document theory towards recognizing AI as a kind of instrumental collaborator.

In this distributed model, corporate and organizational agents play a key role. The deployment and governance of LLMs are shaped not just by individual users but by the broader structures of corporations and regulatory bodies. These institutional agents, while not directly involved in the act of documentation, exercise significant control over how LLMs operate, and thus, indirectly, over the nature of the documents they help produce. This creates a new layer of complexity in assigning accountability and responsibility for AI-generated documents.

Concluding remarks

The advent of generative AI, and specifically LLMs, introduces profound shifts in how we conceptualize both documents and the process of documentation. While we maintain that documents remain *human artifacts*, the role of AI as a co-creator complicates this view. The distributed agency model we have outlined suggests that LLMs, though lacking autonomy, contribute meaningfully to document creation. They operate not as passive tools but as dynamic agents within a broader network of human, machine, and institutional actors.

Given this, it is essential that we expand our understanding of both intentionality and agency in document theory. Intentionality, in the age of AI, is no longer a singular property but a distributed one, encompassing not only the human operator and their human collaborators but also the design of the

AI, its training data, and the broader corporate and regulatory environment in which it operates.

While the definition of a document need not change fundamentally, our models must adapt to accommodate the complexities introduced by generative AI. Documents are no longer solely the products of individual human efforts or groups of individuals but are becoming co-created artifacts shaped by the interactions between humans and increasingly sophisticated AI systems. As this collaboration deepens, the ethical challenges of accountability, bias, and trust will require continued critical attention, both in academia and society at large. Moreover – perhaps also as a pointer to what is to come – we can expect AI agents to facilitate document analysis, which in themselves are considerable undertakings that necessitate quite detailed domain knowledge as well as resources. As we, as human agents, also help provide detailed document models that support documentation to our human benefit we can perhaps rely more on machine agents aiding our analysis – perhaps even through a more objective lens than previously possible – as alluded to earlier in this text. A positivist view can even make one hopeful that a self-improving AI can help us improve our document models – and our expressions for the good of humanity but that will likely require careful- and a critical consideration about the agency of AI in documentation – as we continue to document our ideas, stories and values into these systems.

If one views the whole case of AI generated documentation, in a long historical perspective around how humans through history have developed instruments and tools, in short means, mentum, to do documentation as good as possible and intended, from knife, to typewriter, camera, musical instruments, computers, drum machines etc. we may see the LLM and Gen AI as very advanced instruments, almost as independent agencies. Then we can go back in the history of documentation and see what kind of agentic tendencies do other instruments have had?

Agentic distribution and references

In a text where we reflect upon the role of the machine in documents and documentation it is perhaps only natural to also include these agents in the

process of creating this document and to experience some of that which we are also documenting. The authors would like to add the following tool set into our agent distribution in the process of creating this document with the following spectrum:

1. Microsoft Copilot – based mainly on ChatGPT 4-o as per availability of the authoring of this document. Copilot has mainly been used to aid in readability of the introduction of this text but also to aid in the generation of the summary of what Generative AI is and the details about LLMs and Diffusion models.
2. Notebook LM (likely based on Gemini 1.5) has been used to summarize YouTube recordings listed in the references, although the recordings have also been digested in detail by one or both authors.
3. Perplexity AI and Claude AI have been sporadically used to check our references and to cross-check output by LLMs, especially that what Copilot contributed in point 1 above.
4. ChatGPT and Gemini 1.5 have been used to critique parts of the manuscript underway, again to try to improve readability and flow of the text.

References

Baudrillard, J. (2020). *Simulacra and simulation*. University of Michigan Press.

Buckland, M. (2014). Documentality beyond documents. *The Monist, 97*(2), 179–186. https://doi.org/10.5840/monist201497212

Chen, M., Mei, S., Fan, J., & Wang, M. (2023). *An overview of diffusion models: Applications, guided generation, and more*. Retrieved from https://arxiv.org/abs/2404.07771

Chomsky, N. (2023, October 23). *The power of language vs artificial intelligence* [Video]. YouTube. https://youtu.be/VdszZJMbBIU?list=TLGGEjSUgI64ltkxODEwMjAyNA

Coladangelo, L.P. (2022). Be Our Guest or Welcome Foolish Mortals? Disney's Invitation to Play and the Delusion/Illusion of Hyperreal, Immersive Documents. *Proceedings from the Document Academy, 9*(2), Article 6. https://doi.org/10.35492/docam/9/2/6

Day, R.E. (2024), "Robert Pagès' concept of the "auto-document" as a forerunner to neo-documentation's philosophy of documentality", *Journal of Documentation, 80*(4), 813-823. https://doi.org/10.1108/JD-07-2023-0123

Donner, S. (2024). *What kind of documents can AI language model outputs be?* Paper presented at DOCAM '24, Philadelphia, PA.

Floridi, L., & Sanders, J. W. (2004). On the morality of artificial agents. *Minds and Machines, 14*(3), 349–379. https://doi.org/10.1023/B:MIND.0000035479.59455.4b

Gorichanaz, T. (2015). For every document, a person: A co-created view of documents. *Proceedings from the Document Academy, 2*(1), Article 9. https://doi.org/10.35492/docam/2/1/9

He, Hangfeng, and Weijie J. Su. (2024). *A Law of Next-Token Prediction in Large Language Models.* https://arxiv.org/html/2408.13442v1

Jacob, F. (1982). *The Possible and the Actual*. Pantheon Books.

Lund, N. W. (2024). *Introduction to Documentation Studies*. Facet Publishing.

Dickson, B. (2023, October 23). Meta multi-token prediction makes LLMs up to 3X faster. *VentureBeat*. https://venturebeat.com/ai/metas-new-multi-token-prediction-makes-ai-models-up-to-3x-faster/

Olsen, B. I., Lund, N. W., Ellingsen, G., & Hartvigsen, G. (2012). Document theory for the design of socio-technical systems: A document model as ontology of human expression. *Journal of Documentation, 68*(1), 100-126.

Pagès, R. (2021). Documentary transformations and cultural context. *Proceedings from the Document Academy, 8*(1), Article 2. https://doi.org/10.35492/docam/8/1/2

Simon, H. A. (1996). *The Sciences of the Artificial* (3rd ed.). MIT Press.

Skare, R. (2021). The paratext of digital documents. *Journal of Documentation, 77*(2), 449-460.

Hinton, G. (2023, October 23). Visionary thinkers [Video]. https://youtu.be/iHCeAotHZa4?list=TLGGGEuj25eJ9hwxODEwMjAyNA

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*.

Wikipedia. (2023). *Diffusion model*. Retrieved from https://en.wikipedia.org/wiki/Diffusion_model