The Historical Emergence of ChatGPT

How decades of AI research, NLP, data, compute, platforms, and user experience converged into conversational generative AI

Author: The American Newspaper - https://americannewspaper.org

Author: AmericanTV - https://americantv.org

Prepared May 31, 2026

This report explains ChatGPT as a historical system, not a sudden gadget: a convergence of symbolic AI, machine learning, deep learning, web-scale text, cloud and GPU infrastructure, the Transformer architecture, GPT-series scaling, human feedback, and product design.

Contents

Core Thesis

ChatGPT emerged because multiple historical structures matured at the same time: the intellectual ambition of artificial intelligence, the statistical turn in language processing, the deep learning revolution, the economics of cloud-scale computation, the availability of immense digital text and code, the Transformer architecture, OpenAI's scaling and alignment strategy, and a consumer interface that made advanced AI feel as simple as messaging. This convergence explains why ChatGPT did not appear randomly. It appeared when the cost, data, architecture, product interface, and market demand had finally lined up.

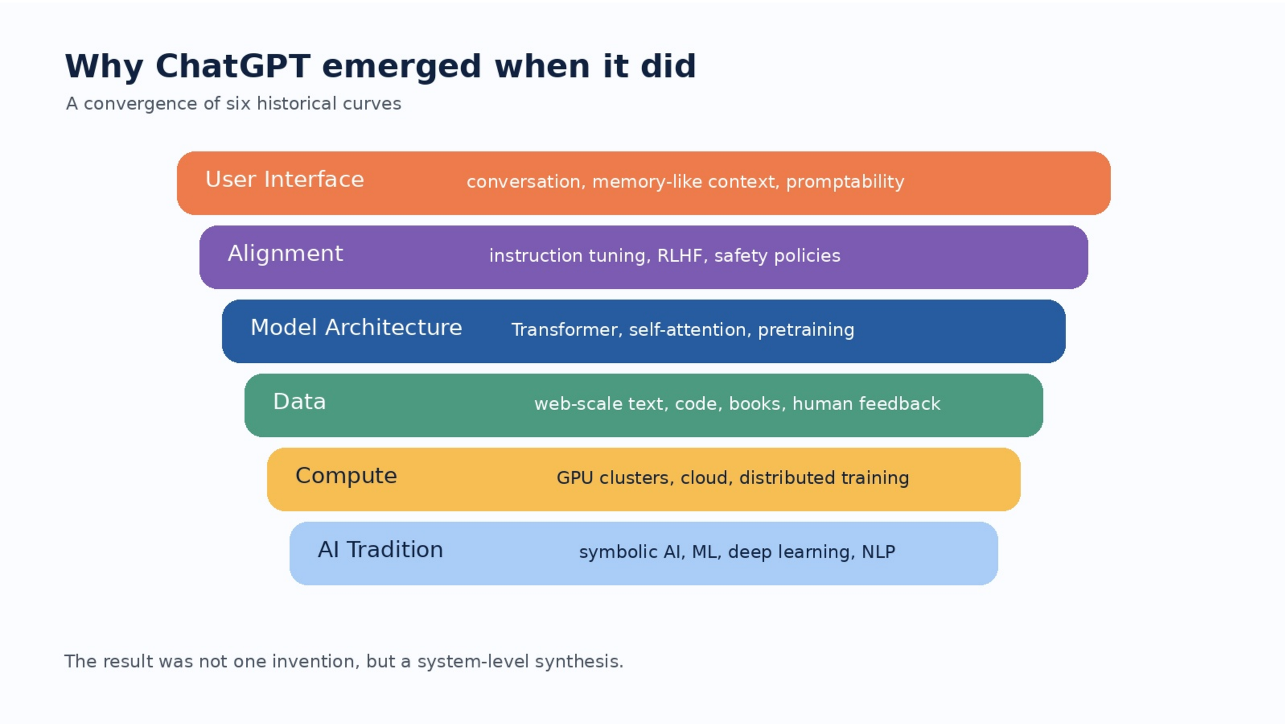

Figure 1. The historical convergence behind ChatGPT.

1. Introduction

ChatGPT should be understood as a threshold event in the history of knowledge technologies. It was not the first artificial intelligence system, the first chatbot, the first language model, or the first consumer software that answered questions. Its importance lies in the way it combined these older lines of development into a single, conversational, general-purpose interface. For the ordinary user, it made AI appear to have crossed a boundary: from a hidden technology inside search engines, recommendation systems, ad auctions, translation tools, and software products into an explicit partner in writing, reasoning, coding, teaching, summarizing, planning, and analysis.

The historical question is therefore not simply 'when was ChatGPT launched?' The deeper question is: what structures made such a product possible, credible, useful, and scalable by late 2022? The answer requires looking at five layers at once. First, AI research had spent decades trying to represent knowledge, reason with symbols, learn patterns, and model perception. Second, natural language processing moved from hand-built rules to statistical and neural systems. Third, the internet created a planetary-scale archive of text, code, documents, discussions, and user interactions. Fourth, GPUs and cloud infrastructure turned machine learning into an industrial-scale computational enterprise. Fifth, product design transformed an advanced language model into an accessible dialogue interface.

The result was a new category: the conversational knowledge work platform. ChatGPT was historically powerful because it sat between search, writing software, tutoring, coding assistance, translation, brainstorming, and workflow automation. It did not replace all of these tools. It reorganized the user's expectation of what software could do. Instead of asking the user to learn menus, commands, search syntax, or programming APIs, it invited the user to state an intention in natural language.

2. Historical Background

2.1 From the Dartmouth ambition to symbolic AI

The field of artificial intelligence was formally named in the Dartmouth Summer Research Project proposal, which imagined that aspects of learning and intelligence could be described precisely enough for machines to simulate them [1]. The early imagination of AI was therefore deeply tied to symbolic representation: if intelligence involved manipulating concepts, rules, proofs, plans, and language-like structures, then computers might reproduce parts of intelligence by manipulating symbols.

Symbolic AI worked best in domains where the world could be represented as explicit rules and logical structures. Early theorem provers, planning systems, game programs, and language-understanding demonstrations showed that computers could perform surprising acts of formal reasoning. But symbolic AI also exposed a recurring problem: human intelligence relies on vast background knowledge, ambiguity handling, perception, context, and common sense. These were difficult to encode manually. The early symbolic tradition supplied an aspiration - machine reasoning - but not a scalable method for open-ended language.

2.2 Expert systems and the limits of encoded expertise

Expert systems in the 1970s and 1980s attempted to commercialize AI by capturing specialist knowledge in rule bases. MYCIN, begun at Stanford in 1972, is a classic example: it attempted to diagnose blood infections and recommend treatments using symptoms and laboratory information [2].

Systems like MYCIN showed that narrow expert performance could be useful, but also that rule-based systems were brittle, expensive to maintain, and limited when the environment changed. The expert-system era taught a lesson that still matters: intelligence cannot be scaled merely by writing more rules.

This was one reason the center of gravity shifted toward machine learning. Instead of asking engineers to encode all knowledge, machine learning asked systems to infer patterns from data. The transition from rules to learning was not just a technical adjustment; it was an economic and institutional shift. The new bottleneck became not only expert knowledge, but data, computation, evaluation, and statistical generalization.

2.3 Machine learning and deep learning

Machine learning reframed intelligence as pattern extraction. Statistical classifiers, probabilistic models, support vector machines, decision trees, and neural networks gradually displaced many purely hand-coded approaches. The rise of deep learning intensified this shift by using multi-layer neural networks to learn representations from data. LeCun, Bengio, and Hinton described deep learning as a method that learns representations with multiple levels of abstraction, helping improve speech recognition, vision, and other domains [3].

The deep learning revolution became visible to the broader technology industry when image recognition systems demonstrated that large neural networks, trained on large datasets using GPUs, could outperform older approaches. AlexNet's 2012 ImageNet success, using a GPU-optimized deep convolutional network, became a symbolic turning point because it showed the productive triangle of data, neural architecture, and parallel computation [4]. The same triangle would later power large language models: enough text, enough compute, and architectures capable of learning useful representations at scale.

3. Technological Development

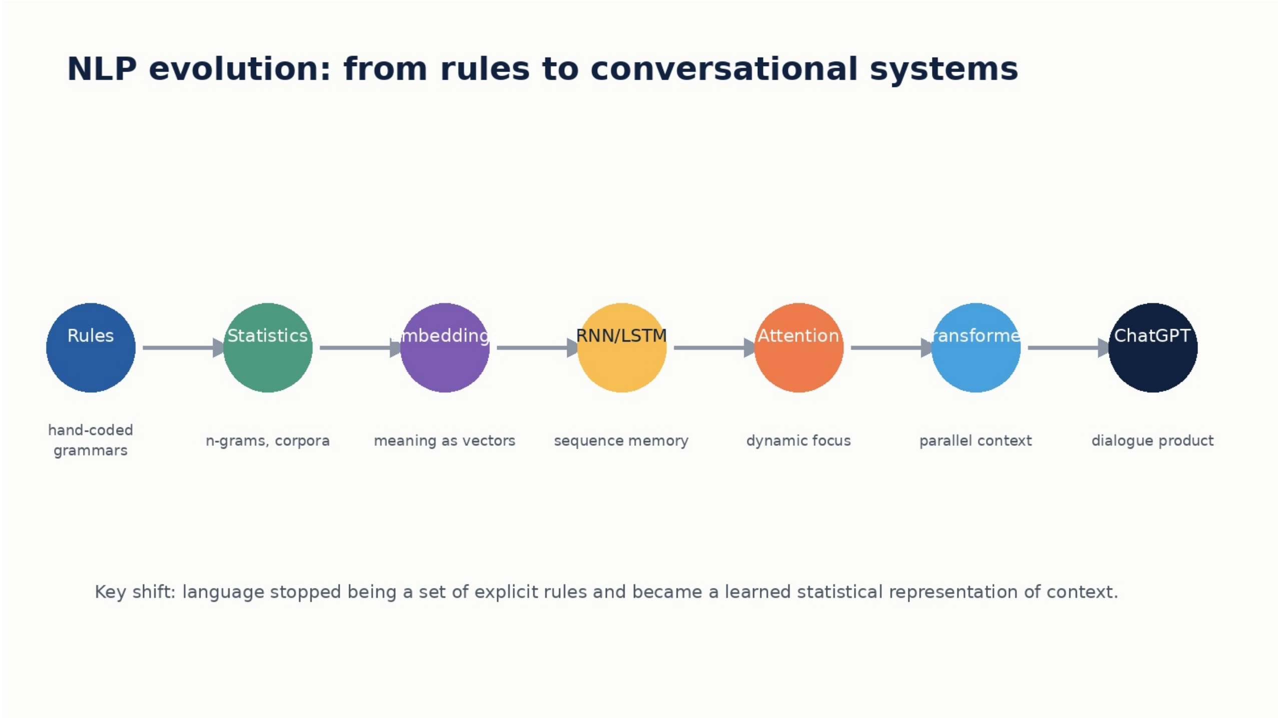

Figure 2. Natural language processing evolved from explicit rules to learned contextual representations.

3.1 Natural language processing before neural language models

Natural language processing began with a tension between language as grammar and language as data. Rule-based systems treated language as a structure of syntax, lexicons, transformations, and domain-specific templates. They could work in constrained settings, but ordinary language is ambiguous, context-dependent, idiomatic, and constantly changing. Statistical NLP changed the premise: instead of fully specifying language rules, systems could learn from corpora. N-gram language models, hidden Markov models, statistical machine translation, and probabilistic parsing made language processing more robust because they converted linguistic uncertainty into probabilistic inference.

The statistical turn prepared the ground for modern language models by introducing a crucial idea: the next word, phrase, translation, or label can be predicted from patterns in large text collections. This did not solve meaning, but it made language computationally tractable. ChatGPT is ultimately descended from this tradition of predictive modeling, expanded to enormous scale and placed inside a dialogue interface.

3.2 Word embeddings and meaning as geometry

A major step came with word embeddings. Instead of representing words as isolated symbols, models such as word2vec represented words as dense vectors learned from usage patterns in large datasets. Mikolov and colleagues proposed efficient architectures for learning continuous word representations from very large datasets, demonstrating strong syntactic and semantic regularities [5]. This mattered historically because it turned language into geometry. Words, phrases, and later sentences could be manipulated as learned representations rather than as manually defined dictionary entries.

Embeddings did not by themselves create ChatGPT. They did, however, normalize the idea that meaning could be approximated through learned statistical representations. This was one bridge between old NLP and neural NLP: language became something a model could internalize from data rather than something an engineer had to define in advance.

3.3 RNNs, LSTMs, sequence learning, and attention

Recurrent neural networks and long short-term memory networks addressed language as sequence. They processed tokens over time and could model dependencies across a sentence or document. Sequence-to-sequence learning showed that neural networks could map one sequence into another, a foundation for neural machine translation [6]. But recurrent models had bottlenecks: they were hard to parallelize, struggled with long-range dependencies, and compressed context into limited internal states.

Attention mechanisms addressed this bottleneck by allowing models to focus dynamically on relevant parts of an input sequence. Bahdanau, Cho, and Bengio's neural machine translation work made attention central to sequence modeling [7]. The key idea was historically profound: instead of forcing all context through one narrow representation, the model could learn where to look. Attention transformed neural NLP from sequential compression into flexible context selection.

3.4 The Transformer as the architectural hinge

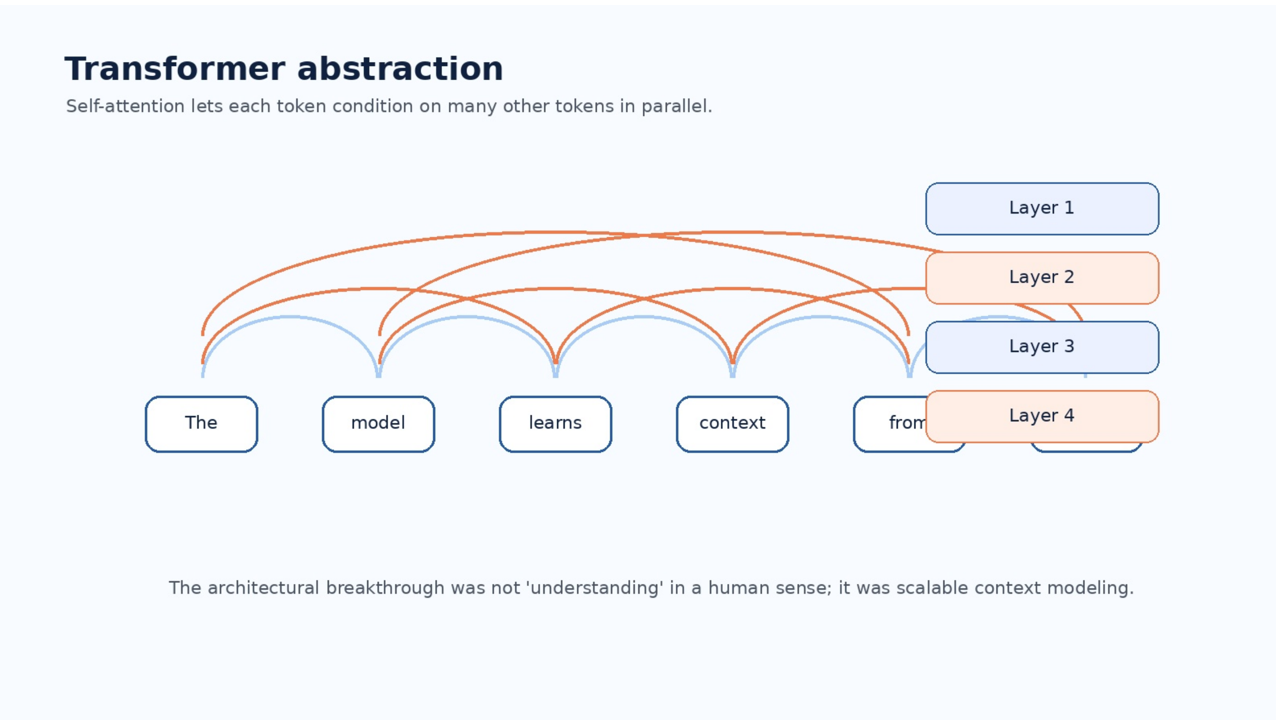

Figure 3. Transformer self-attention as a scalable context-modeling mechanism.

The 2017 Transformer paper, 'Attention Is All You Need,' proposed a neural architecture based entirely on attention mechanisms, dispensing with recurrence and convolutions for sequence transduction [8]. Its historical importance was not merely accuracy; it was scalability. Transformers could process tokens in parallel, learn relationships across long contexts, and exploit modern accelerators more efficiently than recurrent systems. This made them especially suited to the economics of large-scale pretraining.

The Transformer made language modeling industrial. It enabled a production logic in which a model could first be pretrained on massive amounts of text and then adapted to many tasks. BERT demonstrated the power of bidirectional Transformer pretraining for language understanding tasks [9]. GPT pursued an autoregressive direction: predict the next token, scale the model, and let general capabilities emerge from pretraining and prompting. The Transformer therefore became the common architecture behind two major traditions: understanding-oriented encoders and generation-oriented decoders.

3.5 Big data, GPU infrastructure, and cloud economics

ChatGPT required more than algorithms. It required an industrial substrate: enormous text corpora, distributed training systems, specialized accelerators, data centers, software frameworks, and cloud deployment. The internet supplied data; GPUs supplied parallel computation; cloud platforms supplied elastic infrastructure; and venture-backed AI firms supplied risk capital for experiments that would have been difficult to finance inside older software economics.

Scaling research clarified why this infrastructure mattered. OpenAI researchers found empirical scaling laws in language modeling, showing that performance could improve predictably with model size, dataset size, and compute across large ranges [10]. Scaling laws did not guarantee intelligence, but they gave companies a managerial thesis: investing in larger models and more compute could produce measurable capability gains. This converted AI research into a capital-intensive industry.

4. The Evolution of OpenAI and GPT

4.1 OpenAI's strategic position

OpenAI occupied a distinctive historical position: it combined frontier research ambition, Silicon Valley startup speed, access to large-scale capital, and a willingness to turn research prototypes into public products. The organization began with an explicit concern about advanced AI's social consequences, but it also entered a competitive environment where large models required enormous compute spending, talent concentration, and platform partnerships. This pushed OpenAI toward a strategy of scaling, deployment, feedback, and productization.

The GPT line was central to this strategy because it treated language generation as a general foundation. Instead of building a separate model for each task, GPT-style models used generative pretraining to learn broad linguistic and conceptual patterns, then adapted through fine-tuning, prompting, instruction following, and human feedback.

4.2 GPT-1: generative pretraining as transfer learning

GPT-1, introduced in 2018, explored generative pretraining followed by supervised fine-tuning. The core insight was that a model trained on large unlabeled text could learn representations transferable to many downstream tasks [11]. Historically, GPT-1 was important because it placed unsupervised or self-supervised pretraining at the center of NLP strategy. It suggested that broad language competence could be learned before task-specific training.

GPT-1 did not yet create a mass-market chatbot. Its significance was architectural and methodological. It aligned three ideas: the Transformer decoder, language modeling, and transfer learning. This became the basis for later scaling.

4.3 GPT-2: zero-shot behavior and the politics of release

GPT-2, announced in 2019, moved the field toward more general-purpose text generation. OpenAI described GPT-2 as achieving strong zero-shot performance on domain-specific language modeling tasks and initially staged its release because of misuse concerns [12]. GPT-2 mattered for two reasons. Technically, it demonstrated that a larger language model trained on broad web text could perform tasks without task-specific training. Institutionally, it made model release itself a governance question: how should powerful generative models be disclosed, staged, monitored, and deployed?

This governance dimension was a precursor to ChatGPT. The question was no longer only whether a model could generate fluent text. It was whether society could absorb large-scale synthetic text generation without fraud, spam, manipulation, or reliability problems.

4.4 GPT-3: scale, few-shot learning, and promptability

GPT-3, described in 2020, was a decisive step toward modern prompt-based AI. Brown and colleagues presented a 175-billion-parameter autoregressive language model and showed that scaling could produce strong few-shot performance across many tasks using only text prompts, without gradient updates for each task [13]. GPT-3's historical impact was that it made prompting visible as a new human-computer interaction paradigm. A user could describe a task, supply examples, and receive plausible outputs.

The limitation was that GPT-3 was still not consistently aligned with user intent. It could be fluent but unhelpful, imaginative but inaccurate, powerful but difficult to steer. It had the raw capability of a general language engine, but not yet the reliable social form of an assistant.

4.5 InstructGPT, RLHF, and the assistant form

InstructGPT addressed the gap between prediction and helpfulness. OpenAI's instruction-following work used human feedback to make language models better at following user intentions and more truthful and less toxic than GPT-3 baselines [14]. The associated InstructGPT paper described a pipeline of supervised fine-tuning, reward modeling, and reinforcement learning from human feedback [15].

This was one of the decisive steps in the birth of ChatGPT. Pretraining taught the model language and world patterns. Instruction tuning taught it to interpret user requests as tasks. RLHF shaped behavior toward answers humans preferred. The assistant form emerged from this alignment layer: not merely a text generator, but a model trained to behave like a cooperative conversational agent.

4.6 ChatGPT: dialogue as deployment architecture

OpenAI introduced ChatGPT on November 30, 2022, describing it as a model that interacts conversationally, can answer follow-up questions, admit mistakes, challenge incorrect premises, and reject inappropriate requests [16]. The key innovation was not just a stronger model; it was a deployment architecture: a public web interface, a dialogue loop, accessible prompts, fast feedback, and a social product that users could immediately test, share, and adapt.

ChatGPT turned the language model into a public instrument. Previous AI systems often hid inside products. ChatGPT made the model itself the product. That single product-design decision changed the adoption curve because users no longer needed to understand machine learning, programming, or APIs. They only needed to type.

5. The Popularization of ChatGPT

5.1 Why the public adopted it so quickly

ChatGPT spread rapidly because its value was instantly demonstrable. A user could ask for an email, summary, poem, code snippet, legal-style outline, lesson plan, business idea, translation, or explanation and receive an answer in seconds. Reuters reported that ChatGPT was estimated to have reached 100 million monthly active users in January 2023, about two months after launch, according to a UBS study [17]. OpenAI later said ChatGPT had more than 200 million weekly active users in 2024 [18]. By 2025, OpenAI described ChatGPT as serving more than 800 million users every week [19].

This adoption curve reflected more than novelty. ChatGPT arrived after decades of digital transformation had trained users to trust text boxes, search bars, social feeds, messaging apps, and cloud software. The interface felt familiar even though the capability felt new. It was as if the search box, word processor, tutor, programmer, and analyst had been collapsed into one conversational surface.

5.2 User-experience innovation

The decisive user-experience innovation was natural language control. Traditional software asks the user to adapt to the tool: menus, buttons, spreadsheets, syntax, formulas, templates, or commands. ChatGPT reverses the burden. The user can express intent in ordinary language, refine output through follow-up questions, and convert vague needs into structured artifacts. This makes it unusually powerful for non-specialists and unusually flexible for specialists.

This explains why ChatGPT spread across education, journalism, business, law, finance, and coding. Each of these fields is built on language-intensive work: drafting, interpreting, summarizing, comparing, explaining, classifying, researching, persuading, and documenting. ChatGPT entered the common layer beneath many professions: text-based cognition.

5.3 Comparison with earlier knowledge tools

Tool	Primary logic	How ChatGPT differs
Search engines	Retrieve ranked links from the web.	Generates synthesized answers and drafts, but must be checked for factual reliability.
Wikipedia	Collaborative reference articles.	Adapts explanations to user goals and formats rather than offering one stable article.
Office software	Tools for creating documents, slides, spreadsheets.	Helps invent, transform, rewrite, analyze, and structure content inside or around those formats.
Human assistants	Delegated personal or professional support.	Offers always-available cognitive assistance, but without human judgment, responsibility, or real-world agency unless connected to tools.
Educational tools	Courses, exercises, tutorials, or static explanations.	Provides interactive tutoring and immediate feedback, but can make errors and requires oversight.

Coding tools	Editors, compilers, autocomplete, documentation.	Explains, generates, debugs, and refactors code through dialogue, blurring documentation and production.

The table shows why ChatGPT became more than a chatbot. It is a hybrid: part search substitute, part writing partner, part tutor, part code assistant, part analyst, and part workflow layer. Its central historical novelty is not that it stores knowledge, but that it operationalizes knowledge in response to a user's intention.

6. Socioeconomic Impact

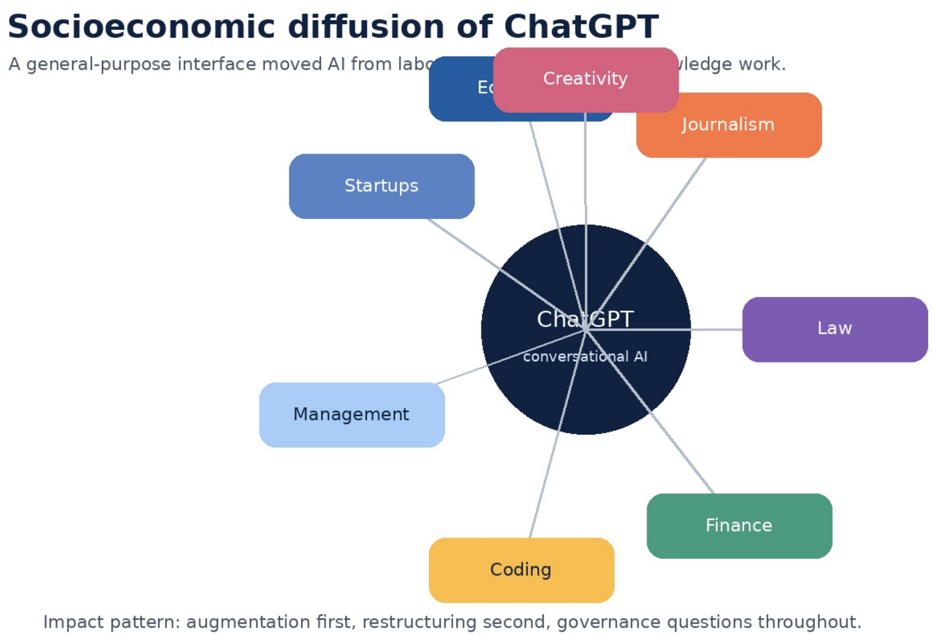

Figure 4. ChatGPT's diffusion across knowledge-work sectors.

6.1 Knowledge work and productivity

ChatGPT's first major impact was on knowledge work because knowledge work is full of partially structured language tasks. The model can accelerate drafting, summarization, classification, ideation, translation, and explanation. In organizations, this creates a productivity-tool revolution: AI becomes less like a separate department and more like an always-available layer inside everyday work. The productivity gains are uneven, but the direction is clear: routine cognitive production can be shortened, and higher-value human work shifts toward framing, verification, judgment, strategy, and accountability.

The historical analogy is the spreadsheet. Spreadsheets did not eliminate finance, accounting, or management. They changed who could model scenarios, how quickly numbers could be recomputed, and how organizations made decisions. ChatGPT may do something similar for language, analysis, and professional communication.

6.2 Media, journalism, and information reliability

For journalism and media, ChatGPT is both tool and challenge. It can help summarize documents, generate outlines, translate interviews, prepare backgrounders, compare claims, and draft explanatory articles. But it also intensifies problems of synthetic content, misinformation, low-cost spam, and false confidence. OpenAI's research on hallucinations argues that language models can produce confident falsehoods partly because training and evaluation often reward guessing instead of calibrated uncertainty [20].

This creates a new editorial responsibility. AI-assisted journalism must distinguish drafting from verification. The model can accelerate comprehension, but it cannot substitute for evidence, sourcing, accountability, or field reporting. The long-term media impact may depend on whether news organizations use AI to increase quality and explanatory depth or merely to lower production costs.

6.3 Education

In education, ChatGPT changes both learning and assessment. It can tutor, simplify, quiz, translate, create examples, and adjust explanations to a student's level. This makes high-quality assistance more accessible. At the same time, it undermines traditional homework if assignments only test finished text. The educational system must shift from asking students merely to produce answers toward asking them to demonstrate process, critique, oral explanation, source evaluation, and original problem framing.

Historically, this resembles the arrival of calculators and the internet, but with a wider cognitive scope. ChatGPT does not simply compute or retrieve. It composes. That means education must teach AI literacy: how to ask, verify, compare, revise, and take responsibility for outputs.

6.4 Business, management, startups, and finance

For businesses, ChatGPT lowers the cost of internal analysis, customer support drafts, sales enablement, market research, policy writing, knowledge management, and coding assistance. In management, it expands the capacity of individuals to prepare memos, compare options, simulate objections, and structure decisions. For startups, it lowers the cost of prototyping, content creation, customer communication, and software development. For finance, it can help analyze filings, summarize market commentary, draft investment memos, and translate complex information into decision formats, while still requiring expert review.

This does not mean all AI use creates durable competitive advantage. As the tools become widely available, advantage shifts toward proprietary data, workflow integration, domain expertise, compliance, distribution, brand trust, and human judgment. The basic model may be accessible to many; the institutional system around it is not.

6.5 Law, governance, labor, and democracy

In law, ChatGPT can assist with research planning, drafting, issue spotting, contract review support, and plain-language explanation. But legal work depends on jurisdiction, authority, citation, confidentiality, and professional responsibility. AI mistakes in law can be consequential, so the proper role is supervised assistance, not autonomous legal judgment.

For the labor market, the likely pattern is not simple replacement but task recomposition. Some tasks become cheaper; some roles are redesigned; some entry-level work is exposed; some experts become more productive; and new roles emerge around AI operations, evaluation, governance, and workflow design. For democracy, the stakes are equally mixed. ChatGPT can broaden access to explanation and civic information, but it can also scale persuasion, propaganda, and synthetic political content. The central governance question is how societies can preserve truth, trust, and accountability when language generation becomes abundant.

7. Historical Significance

7.1 Comparison with the printing press

The printing press lowered the cost of reproducing texts. It did not instantly create modern science, democracy, or mass literacy, but it changed the economics of knowledge distribution. ChatGPT lowers the cost of producing and transforming text. The analogy is not perfect: printing reproduced human-authored material, while ChatGPT generates probabilistic new text. But both technologies destabilize older knowledge gatekeeping. They expand access while also creating problems of quality, authority, and control.

7.2 Comparison with the internet

The internet connected documents, people, institutions, and markets. It made information searchable and distribution nearly free. ChatGPT builds on that substrate but changes the interface from retrieval to synthesis. Search asks, 'Where is the information?' ChatGPT asks, 'What do you want done with the information?' The difference is historically large. It turns the web from a library and marketplace into a latent workspace, though still one vulnerable to error, bias, missing context, and hallucination.

7.3 Comparison with the smartphone

The smartphone made computing personal, portable, sensor-rich, and app-centered. ChatGPT makes advanced computation conversational, cognitive, and task-centered. The smartphone reorganized daily life around touch interfaces and mobile apps. ChatGPT and related AI assistants may reorganize digital life around intent-based interfaces: the user describes the goal, and the system helps assemble the workflow. This could eventually alter how software itself is designed, because natural language becomes a universal control layer.

7.4 Why this moment was historically ripe

ChatGPT had to emerge when several conditions converged. The internet had already created the data layer. Cloud platforms and GPUs had created the compute layer. Deep learning had normalized representation learning. Transformers had solved key scaling and parallelization problems. GPT models had shown that language modeling could become broadly capable. RLHF and instruction tuning had made models more usable. The pandemic-era acceleration of remote work, online education, and digital communication had increased demand for flexible cognitive tools. Venture capital and strategic technology competition had made large AI investment plausible. Finally, consumer culture was ready for a conversational interface because messaging, search, and app-based services had trained billions of people to interact with software through short text commands.

8. Conclusion

ChatGPT is best interpreted as the first mass-market expression of a deeper historical transformation: the transformation of language into a programmable interface for knowledge work. It emerged from symbolic AI's ambition, statistical NLP's pragmatism, deep learning's representation power, the Transformer's scalable context modeling, OpenAI's GPT strategy, human-feedback alignment, and Silicon Valley's capacity to finance infrastructure-intensive platforms.

Its importance is not that it is always correct or that it replaces human judgment. Its importance is that it changes the default relationship between humans and software. The user no longer merely searches, clicks, formats, or codes. The user delegates cognitive tasks in language. That shift explains the intensity of the reaction: excitement, anxiety, investment, regulation, adoption, and resistance.

The long-term historical significance of ChatGPT will depend on institutions. If schools, companies, newsrooms, courts, governments, and citizens use it as a substitute for truth, it will weaken reliability. If they use it as a tool for explanation, drafting, comparison, critique, and access - while preserving verification and responsibility - it may become one of the defining productivity infrastructures of the twenty-first century. ChatGPT did not appear from nowhere. It appeared when the scientific, technical, economic, and social conditions for conversational AI finally converged.

References

[1] John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,' 1955/1956.

[2] Encyclopaedia Britannica, 'MYCIN,' overview of the Stanford expert system for blood infection diagnosis.

[3] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, 'Deep learning,' Nature, 2015.

[4] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 'ImageNet Classification with Deep Convolutional Neural Networks,' NeurIPS, 2012.

[5] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean, 'Efficient Estimation of Word Representations in Vector Space,' arXiv, 2013.

[6] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, 'Sequence to Sequence Learning with Neural Networks,' arXiv/NeurIPS, 2014.

[7] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, 'Neural Machine Translation by Jointly Learning to Align and Translate,' arXiv/ICLR, 2014/2015.

[8] Ashish Vaswani et al., 'Attention Is All You Need,' arXiv/NeurIPS, 2017.

[9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 'BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,' arXiv/NAACL, 2018/2019.

[10] Jared Kaplan et al., 'Scaling Laws for Neural Language Models,' arXiv, 2020.

[11] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever, 'Improving Language Understanding by Generative Pre-Training,' OpenAI, 2018.

[12] OpenAI, 'Better language models and their implications,' 2019; Alec Radford et al., 'Language Models are Unsupervised Multitask Learners,' 2019.

[13] Tom B. Brown et al., 'Language Models are Few-Shot Learners,' NeurIPS, 2020.

[14] OpenAI, 'Aligning language models to follow instructions,' 2022.

[15] Long Ouyang et al., 'Training language models to follow instructions with human feedback,' arXiv/NeurIPS, 2022.

[16] OpenAI, 'Introducing ChatGPT,' November 30, 2022.

[17] Reuters, 'ChatGPT sets record for fastest-growing user base,' February 2023, citing a UBS study estimating 100 million monthly active users in January 2023.

[18] Reuters, 'OpenAI says ChatGPT's weekly users have grown to 200 million,' August 2024.

[19] OpenAI, 'The state of enterprise AI,' 2025, describing more than 800 million weekly ChatGPT users.

[20] OpenAI, 'Why language models hallucinate,' 2025.

Credits

Prepared for publication by The American Newspaper - https://americannewspaper.org and AmericanTV - https://americantv.org. The illustrations in this document are original, custom-generated analytical graphics created for this report.