

B. O. KUZIKOV

Candidate of Technical Sciences, Associate Professor,
Associate Professor at the Department of Computer Science
Sumy State University
ORCID: 0000-0002-9511-5665

S. R. SHOVKOPLIAS

Postgraduate Student at the Department of Computer Science
Sumy State University
ORCID: 0000-0003-1837-0213

S. P. SHAPOVALOV

Candidate of Physical and Mathematical Sciences, Associate Professor,
Senior Lecturer at the Department of Computer Science
Sumy State University
ORCID: 0000-0002-4936-9461

ITERATIVE FORMALIZATION OF EXPERT KNOWLEDGE INTO SYSTEM INSTRUCTIONS FOR LARGE LANGUAGE MODELS

*This paper addresses the problem of logical inference degradation in large language models (LLMs) during semantic analysis of strategic documents exceeding 100 pages in length. The study focuses on developing a framework for the iterative formalization of tacit expert knowledge into system instructions to automate the auditing of community development plans for compliance with gender-sensitive recovery standards. **Methodology.** A multi-level logical inference approach, structured as an Orchestrator–Executor architecture, was implemented using models of the gemini-2.5-flash class. The mathematical model of the system is represented as a tuple of operators comprising: an Initializer–Extractor for reverse engineering of expert reasoning; an Executor; a QA-Auditor for hallucination detection; a Judge; and a Trainer responsible for constructing an explanatory model of errors. Optimization is realized as an iterative cycle incorporating a Human-in-the-Loop (HITL) mechanism to prevent semantic drift and the loss of domain-specific terminology. To eliminate anachronisms, a semantic version control mechanism was introduced that adapts analytical criteria to the historical context of the document under review. **Results.** Verification was conducted through blind testing on strategic plans of 15 territorial communities. The results confirmed the high generalization capacity of the optimized prompt, achieving with a validation level approximating that of professional human assessment on a 0–10 point scale. The application of the Orchestrator–Executor architecture enabled a substantial reduction in operational costs while preserving the accuracy of fact extraction. **Novelty.** The conceptual contribution lies in reconceptualizing knowledge distillation as a process of "crystallizing" expert heuristics into the semantic space of system prompts, without requiring neural network fine-tuning. **Practical Significance.** The practical relevance of the research is demonstrated within the framework of the GenPlan project, where the developed toolkit provides screening of strategic documentation in support of gender-sensitive community recovery planning.*

Key words: large language models; LLM; knowledge formalization; prompt optimization; automated document auditing; iterative instruction tuning; gender-sensitive recovery.

Б. О. КУЗІКОВ

кандидат технічних наук, доцент,
доцент кафедри комп'ютерних наук
Сумський державний університет
ORCID: 0000-0002-9511-5665

С. Р. ШОВКОПЛЯС

аспірант кафедри комп'ютерних наук
Сумський державний університет
ORCID: 0000-0003-1837-0213



С. П. ШАПОВАЛОВ

кандидат фізико-математичних наук, доцент,
старший викладач кафедри комп'ютерних наук
Сумський державний університет
ORCID: 0000-0002-4936-9461

ІТЕРАТИВНА ФОРМАЛІЗАЦІЯ ЕКСПЕРТНИХ ЗНАНЬ У СИСТЕМНІ ІНСТРУКЦІЇ ВЕЛИКИХ МОВНИХ МОДЕЛЕЙ

Стаття присвячена розв'язанню проблеми деградації логічного виводу великих мовних моделей (LLM) при семантичному аналізі великих документів. Дослідження фокусується на розробці фреймворку ітеративної формалізації неявних експертних знань (*tacit knowledge*) у системні інструкції для автоматизації аудиту планів розвитку громад на відповідність стандартам гендерно-чутливого відновлення. **Методологія.** Застосовано підхід багаторівневого логічного виводу («Оркестратор–Виконавець») на базі моделей класу *gemini-2.5-flash*. Математична модель системи представлена кортежем операторів, що включає Ініціалізатор-екстрактор для реверс-інжинірингу експертного мислення, Виконавця, QA-Аудитора для детекції галюцинацій, Суддю та Тренера, який формує пояснювальну модель помилок. Оптимізацію реалізовано як ітеративний цикл із механізмом *Hitap-in-the-loop* (HITL) для запобігання семантичному дрейфу та втраті фахової термінології. Для усунення анахронізмів впроваджено механізм семантичного контролю версій, що адаптує критерії аналізу до історичного контексту документа. **Результати.** Верифікацію проведено шляхом «сліпого тестування» на стратегічних планах 15 територіальних громад. Результати підтвердили високу узагальнювальну здатність оптимізованого промпту, забезпечивши рівень валідації, наближений до професійної людської оцінки за шкалою 0–10 балів. Використання архітектури «Оркестратор–Виконавець» дозволило суттєво знизити операційні витрати зі збереженням точності екстракції фактів. Новизна полягає у переосмисленні дистіляції знань як процесу «кристалізації» експертних евристик у семантичний простір системних промптів без донавчання нейромережі. Практична значущість реалізована в межах проекту «GenPlan», де інструментарій забезпечує скринінг стратегічної документації для гендерно-чутливого планування відновлення громад.

Ключові слова: великі мовні моделі; LLM; формалізація знань; оптимізація промптів; автоматизований аудит документів; ітеративна оптимізація інструкцій; гендерно-чутливе відновлення.

Introduction

Large language models (LLMs) are reshaping the field of intelligent data analysis. Advances in the field have driven a shift from deploying neural networks as auxiliary tools for text generation towards their integration into complex analytical pipelines as autonomous agents [1]. Such systems can address complex tasks that previously required specialist human involvement. Nevertheless, the practical deployment of LLMs in professional settings has revealed limitations that cannot be overcome simply by scaling computational capacity.

One particularly challenging scenario is the analysis of lengthy strategic documents – materials that, within the domain of public administration, routinely run to between 100 and 120 pages. Although contemporary models operate with context windows exceeding one million tokens, their actual processing capacity remains constrained. Research has identified the phenomenon of ‘illusory reasoning’: beyond a certain level of task complexity, model inference capabilities deteriorate sharply. This ‘accuracy collapse’ occurs even when a sufficient token budget is available, rendering traditional approaches – most notably few-shot prompting – ineffective, as models lose structural coherence when processing large volumes of data.

This challenge arose within the ‘Integrated Gender-Sensitive Community Recovery’ project, where the central task is auditing strategic plans for compliance with national standards and international Sustainable Development Goals – in particular, SDG 5 and Order No. 86 of the Ministry of Social Policy [2]. The conventional approach to such assessment faces a significant scalability problem: a shortage of qualified specialists, combined with the considerable cognitive burden of manually reviewing hundreds of pages of bureaucratic text, results in mounting ‘expert noise’. At the same time, direct LLM-based automation frequently produces unsatisfactory outcomes, owing to the models’ propensity for hallucinations and their tendency to disregard anachronisms – both of which are inadmissible in professional evaluation practice.

This study responds to the need for hybrid systems that combine the flexibility of language models with the rigorous logic of formalised audit criteria. Rather than relying on static prompt construction, this work develops an approach to the iterative optimisation of system instructions, enabling the distillation of tacit expert knowledge into a functional prompt – effectively encoding expert heuristics into the model’s context.

The aim of the study is to develop and experimentally validate a framework for adaptive knowledge distillation that automates the initial stages of document analysis. Such tooling is intended to substantially reduce the workload of expert auditors during the drafting and screening of strategic plans, whilst achieving a level of validation close to professional human assessment – at significantly lower operational cost.

Theoretical background and review of recent research

The development of methods for governing the behaviour of large language models has evolved from basic prompt formulation (zero-shot) to sophisticated engineering frameworks. Research [3] has confirmed the efficacy of the Chain-of-Thought (CoT) method, which encourages the model to make its intermediate reasoning steps explicit. For tasks operating within strict normative constraints, however, CoT proves insufficient in the absence of an external control loop. Recent studies [4] indicate that even advanced large reasoning models (LRMs) experience an ‘accuracy collapse’ as task complexity increases – a finding that makes the shift to a prompt programming paradigm timely and necessary. Frameworks such as DSPy treat the instruction not as a static piece of text but as a function parameter [5], subject to optimisation via mechanisms such as OPRO (Optimisation by PROMpting) or MIPRO [6].

In the context of in-context learning, knowledge distillation takes on a different meaning: rather than updating neural network weights according to Hinton’s method [7], expert experience is crystallised into a highly effective system prompt. This allows tacit knowledge to be conveyed to the model – intuition, an understanding of professional context, and the appropriate register of critique – without the prohibitive cost of fine-tuning. A further important consideration is the incorporation of a Human-in-the-Loop (HITL) mechanism [8], wherein a human operator acts as a ‘Meta-Judge’. This safeguard prevents ‘semantic drift’ and reward hacking, which arise when fully automated systems (AutoML) optimise formal metrics at the expense of analytical substance.

Particular attention must be given to the limitations of few-shot prompting when applied to documents with heterogeneous structures. Providing an example analysis for one structural type – a financial report, for instance – can disorient the model when it encounters a legal contract, compelling it to search for non-existent parallels and, consequently, to generate hallucinations. Furthermore, an excessive number of examples consumes context window resources, which is especially critical when processing documents in excess of 100 pages. A comparative analysis of prompting strategies is presented in Table 1.

Table 1

Comparative Analysis of Prompting Strategies

Prompt Engineering Strategy	Operating Principle	Effectiveness for Heterogeneous Structures	Risk of Hallucination	Optimal Domain of Application
Zero-Shot	Direct instruction with no examples.	Low (model-dependent)	Very high	Simple classification, translation
Few-Shot	Provision of N input-output examples	Critically low (imposes patterns)	Moderate	Entity extraction from uniform reports
Chain of Thought (CoT)	Step-by-step generation of reasoning	High (adapts to document structure)	Low	Identification of complex relationships in text
Self-Consistency	Majority selection across multiple CoT runs	Very high (mitigates individual errors)	Minimal	Extraction of quantitative indicators
Tree of Thoughts (ToT)	Proposal, evaluation, and synthesis of hypotheses	Absolute (expert-level performance)	Minimal	Multi-criteria analysis

The effectiveness of any given strategy is directly contingent on the configuration of the architectural parameters governing access to the LLM. For tasks requiring rigorous analysis in which the precision of fact extraction is paramount, the Temperature parameter should be reduced to its minimum (0.0–0.1); combined with a constrained Top-P value, this renders the model’s output deterministic. In this configuration, the system message imposes strict constraints and imperatives – for example, a prohibition on drawing upon external data – whilst the user message serves as a container for dynamic instructions and the text under analysis.

Task formalization

The system is formalised as a mapping from a fuzzy space of semantic features to a well-defined set of analytical decisions. The mathematical model is expressed as the following tuple:

$$M = D, E, P, R, S, H; \Phi_1, \Psi, \Theta, \mathcal{J}, \Delta, \Omega$$

where: $D = \{D_{lab}, D_{unlab}\}$ – the set of labelled and unlabelled documents; E – the space of reference expert conclusions (the "gold standard"), containing implicit evaluation criteria; P – the semantic space of system prompts (instructions); R – the set of AI-generated reports; S – the state vector of error memory (regression basis); H – the set of operator control inputs (hints) provided by the human operator. Functional operators of the system: Φ_1 (Initialiser-Extractor) – performs reverse engineering of expert reasoning on the basis of the a priori corpus of texts E , extracting implicit criteria and forming the initial prompt $p^{(1)} \in P$; Ψ (Executor) – maps the pair $\{p^{(k)}, D\}$ to a report $r \in R$. Implemented using Fast Model-class models to optimise operational costs; Θ (QA Auditor) – validates outputs on unlabelled data D_{unlab} , checking for hallucinations and anachronisms. Should the hallucination rate exceed the permissible threshold θ , a "Penalty Loop" is activated, thereby preventing overfitting of the system; \mathcal{J} (Judge) – computes the divergence between r and $e \in E$,

classifying discrepancies as omissions (K_{fn}) or fabrications (K_{fp}); Δ (Trainer) – aggregates errors in light of the history $S^{(k)}$ and constructs an explanatory model of the causes of non-conformance; Ω (Optimiser) – effects the transition to a new prompt state $p^{(k+1)}$, maximising the integral conformance metric $J(p)$.

The objective function $J(p)$ is defined as:

$$J(p) = w_1 \cdot \text{Sim}(R, E) + w_2 \cdot \left(1 - \frac{K_{fn} + K_{fp}}{N} \right)$$

where $\text{Sim}(R, E)$ denotes the semantic relevance score determined by means of gap analysis, and w_i are weighting coefficients that set the priority of either analytical precision or completeness.

This architecture enables the system to learn autonomously from the unstructured outputs of an expert, progressively formalising them into stable analytical instructions.

Research methodology

Functional Model of Differential Optimisation and Feedback

The central element of the system is the Trainer operator (Δ), which acts as a link between the automated assessment performed by the Judge (\mathcal{J}) and the strategic decisions of the human operator. Rather than simply aggregating scores, ” produces an explanatory error model (Explainable AI), enabling decisions to be taken at the level of architectural changes to the prompt – rather than through the piecemeal correction of individual outputs.

The Trainer’s operating logic rests on three core analytical procedures.

- The *regression audit procedure* involves a comparative analysis of the current error matrix against the state vector $S^{(k)}$, enabling the identification of "regression" instances – situations in which the correction of specific defects triggers the re-emergence of previously resolved errors – thereby guaranteeing the temporal stability of the system.

- The *error consensus detection mechanism* is designed to identify persistent patterns across the corpus of processed documents. Discrepancies – in particular, false negatives – that appear in more than 80% of generated reports are classified as a fundamental systemic deficiency in the instructions, rather than as random anomalies attributable to a specific model.

- *Root cause generation* is accomplished by establishing deterministic causal relationships between the architecture of the system prompt and the identified deviations. Where systematic hallucinations are detected, for instance, this mechanism enables a critical examination of the instruction hierarchy for imbalances between demands for detail and the imperative of factual accuracy.

By separating the functions of the Executor (Ψ), the Judge (\mathcal{J}), and the Trainer (Δ), the system realises the principle of hierarchical learning. This makes it possible to delegate routine data processing to "lightweight" models, which radically reduces the per-unit cost of analysis and minimises the system's carbon footprint through the use of energy-efficient, compact models.

In operational mode, the optimal prompt p^* is fixed upon discovery, ensuring strong generalisation performance on documents outside the training set whilst preserving expert-level validation standards.

Algorithmic Implementation of the Iterative Optimisation Cycle

The practical implementation of the proposed method (Fig. 1) is based on a software suite developed in the Ruby programming language. The approach involves a hierarchical allocation of computational resources between two model classes: a high-capacity executor model with an extended context window for the direct processing of large document corpora, and an intelligent analyst model that serves as analytical process orchestrator.

The iterative training cycle is implemented as a closed loop, commencing with the knowledge base initialisation stage, during which the analyst model performs reverse engineering of reference expert conclusions to extract implicit evaluation criteria. The resulting initial system prompt is then applied in batch mode to the training document set via the executor model. This is followed by a comparative evaluation stage, in which the analytical core cross-references the generated outputs against the reference values and classifies any discrepancies according to three categories: semantic hallucinations, logical deficits, and chronological aberrations.

A critical component of the cycle is the Human-in-the-Loop mechanism, which involves presenting the error analysis report to a human operator who makes strategic decisions regarding the direction of the system’s further development. On the basis of this feedback, differential optimisation is performed: an automated optimiser rewrites the system prompt to address identified deficiencies, whilst mandatory regression testing is conducted in parallel.

Particular attention within the algorithm is devoted to resolving the problem of evaluating documents from prior periods, which language models frequently analyse through the lens of contemporary wartime and socio-political realities. To counteract this effect, a semantic version control block was integrated into the system instruction architecture, which compels the model to identify the temporal attributes of a document prior to commencing analysis. This enables the system to switch between Mode A – an extended audit for compliance with wartime regulatory requirements – and Mode B – an evaluation of documents that accounts for their historical context, wherein outdated data is not treated as erroneous.

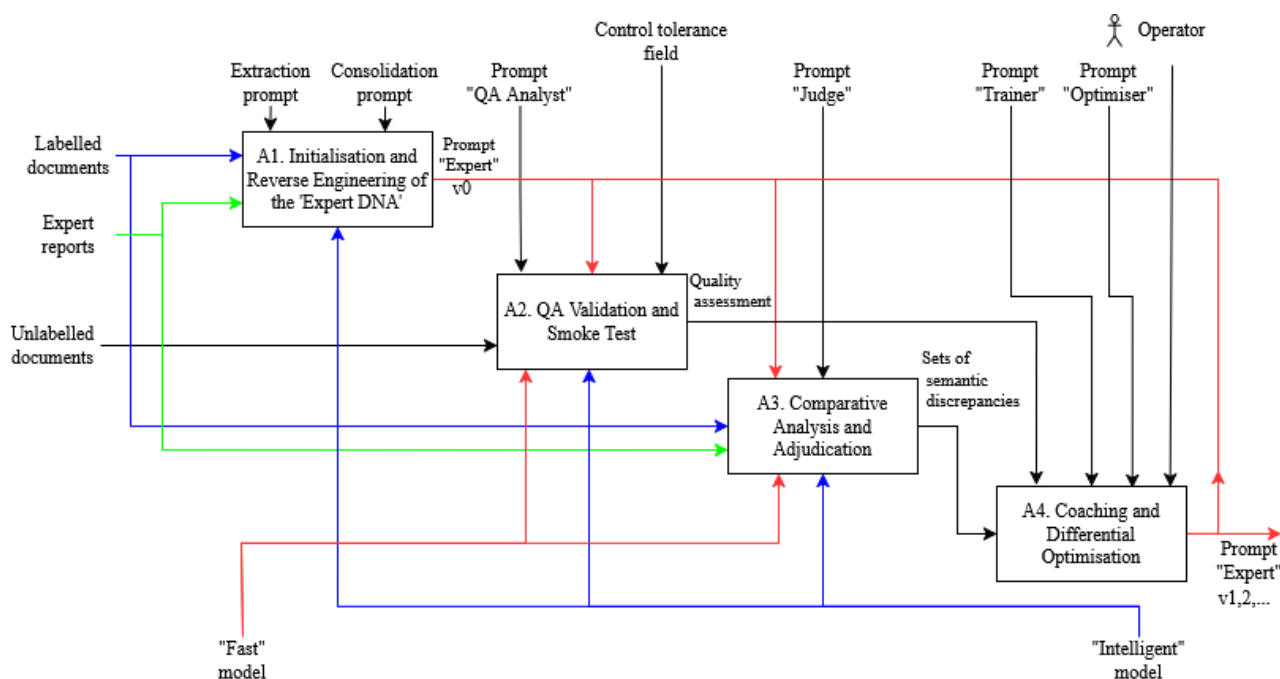


Figure 1. IDEF0 model of the iterative system prompt optimisation cycle

Verification Procedure and Assessment of Generalisation Capacity

To confirm the reliability of the optimised prompt and rule out the risk of overfitting to a limited dataset, a blind testing procedure was implemented. Verification was conducted on a control sample comprising the strategic plans of 15 territorial communities that had not been involved in any stage of the iterative optimisation process. The verification methodology is based on autonomous internal consistency auditing and incorporates the deterministic generation of reports by the executor model at a temperature value of zero. The verification process involves QA analysis across several dimensions:

- The system performs *factual hallucination detection* by identifying references to non-existent legislative acts or fabricated statistical indicators.
- The analytical core identifies *logical deficits* in cases where the model produces unwarrantedly positive conclusions in the absence of supporting factual data in the source document.
- The model carries out *chronological validation* to ensure that temporal context is respected and that no anachronisms are present in the conclusions.
- *Stylistic verification* is conducted to assess the objectivity of the text and to eliminate instances of artificial empathy or emotional colouring.

The outcome of the verification process is a composite quality score ranging from 0 to 10. A maximum score denotes full conformity of the report with the content of the original document and strict adherence to an analytical form of presentation. This procedure enables an objective assessment of the system’s capacity to produce reliable analyses of new documents without requiring additional hand-crafted reference reports.

Experiment and results

Research Object and Experimental Conditions

The study drew on strategic documents from four Ukrainian territorial communities of varying types, including a regional centre, a district-level town, and a border community. The adaptive knowledge distillation process encompassed five complete cycles of iterative instruction optimisation. In the course of these trials, a critical vulnerability of autonomous systems was identified: during the third iteration, in the absence of direct human intervention, semantic drift was observed. The model began to replace domain-specific terminology, substituting distinctly Ukrainian administrative concepts – such as *starosta* district and *hromada* – with generic anglicisms: ‘*municipality*’ and ‘*district*’.

To stabilise the lexical composition of the reports in subsequent cycles, the Human-in-the-Loop mechanism was employed, enabling the introduction of a hard negative constraint prohibiting the use of foreign-language terminology. This preserved the conceptual vocabulary characteristic of Ukraine’s normative framework, and confirmed that human oversight is necessary in the knowledge distillation process when operating within highly specialised domains.

Verification and Analysis of Generalisation Capacity

To confirm the reliability of the optimised instruction and rule out the risk of overfitting on a limited dataset, a blind testing procedure was implemented. The control sample comprised strategic plans from 15 territorial communities

sourced from open repositories, none of which had been involved in the iterative search for the optimal prompt. The verification methodology was based on autonomous internal consistency auditing and incorporated several stages of independent evaluation, including expert involvement to provide external validation and support the results.

The quality assessment process for the generated reports followed a clearly defined algorithm.

- The *executor model* produced analytical reports for the validation sample exclusively in deterministic mode, with output temperature set to zero, to guarantee the reproducibility of results.
- The *analyst model* assumed the role of a rigorous auditor, conducting a multi-factor QA revision analysis to detect factual hallucinations, logical deficits, and chronological aberrations.
- The *composite quality score* (QA score, 0–10) enabled an objective assessment of the degree to which each report conformed to the actual content of the original document and met the requirements of a strictly analytical form of presentation.

The results confirmed that the blind testing stage is essential for identifying hidden prompt deficiencies, which tend to be masked during operation with the primary training sample. The domain experts involved in the study confirmed that the final AI-generated reports align, in both content and depth of analysis, with the original expert conclusions. This demonstrates the system's strong capacity to generalise expert approaches to new communities without additional hand-crafted reference examples. The method has been implemented within the Integrated Gender-Sensitive Community Recovery project and is available for testing on a dedicated domain.

Conclusions

This study developed and validated a method for the adaptive distillation of expert knowledge into large language models for automated strategic document auditing. The proposed framework of iterative system instruction optimisation overcomes key limitations of contemporary Transformer architectures – most notably, the degradation of logical inference when processing large volumes of heterogeneous data.

The principal contribution of this work is the reframing of knowledge distillation as a process of crystallising tacit expert knowledge directly into the semantic space of a system prompt. The implemented Orchestrator–Executor architecture, built upon the gemini-2.5-flash and gemini-2.5-flash-lite models, demonstrated high analytical accuracy whilst substantially reducing operational costs. The architecture enabled the separation of routine text processing from high-level logical analysis – a distinction important for the scalability of digital governance systems.

Experimental validation confirmed that integrating a Human-in-the-Loop mechanism is essential for preventing semantic drift and preserving domain-specific terminology. The introduction of a semantic version control block successfully mitigated the problem of anachronisms, ensuring the correctness of document evaluation across different time periods whilst accounting for the relevant historical and normative context. The blind testing results demonstrated strong generalisation capacity of the optimised prompt p^* , whose performance on previously unseen documents closely approximates that of professional expert assessment.

The practical relevance of the research has been demonstrated within the GenPlan (Gender responsive development and recovery planning for communities) project (<https://genplan.csgard.sumdu.edu.ua/>), where the developed tooling provides initial screening and scenario planning for the gender-sensitive recovery of communities. Broader implementation of such hybrid intelligent systems within local self-governance bodies will contribute to greater transparency in administrative decision-making and a substantial reduction in the workload of expert auditors in the drafting of strategic documents.

Acknowledgements. The information system under development was created within the framework of the project "Integrated Gender-Sensitive Community Recovery: Data Analysis and Scenario Planning Tools", carried out by the Civil Society Organisation "Center for Public Initiatives "Intellect of Sumy Region"". The authors express their gratitude to experts Sofiia BUDNYK, Tetiana HREBENYK, Liudmyla KALINICHENKO, and Svitlana PIVEN for their fruitful collaboration in validating the results obtained.

Bibliography

1. Новаковський А., Яловега І. Визначення ролей агентів великих мовних моделей (LLM) у моделі дизайн-процесу. *Сучасний стан наукових досліджень та технологій в промисловості*. 2025. № 3 (33). С. 58–72. DOI: 10.30837/2522-9818.2025.3.058.
2. Про затвердження Методичних рекомендацій щодо оцінювання гендерного впливу галузевих реформ : Наказ Міністерства соціальної політики України від 07.02.2020 № 86.
3. Wei J., Wang X., Schuurmans D. et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*. 2022. Vol. 35. P. 24824–24837.
4. Shojaee P., Mirzadeh I., Alizadeh K. et al. The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity. Apple, 2025. 23 p. DOI: 10.48550/arXiv.2506.06941
5. Khattab O. et al. DSPy: Compiling Declarative Language Model Calls into Self-Improving Pipelines. *arXiv preprint*. 2024. DOI:10.48550/arXiv.2310.03714

6. Yang C., Wang X., Lu Y. et al. Large language models as optimizers. *arXiv preprint*. 2024. DOI:10.48550/arXiv.2309.03409
7. Hinton G., Vinyals O., Dean J. Distilling the knowledge in a neural network. *arXiv preprint*. 2015. DOI:10.48550/arXiv.1503.02531
8. Wu X., Xiao L., Sun Y., Zhang J., Ma T., He L. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems*. 2022. Vol. 135. DOI: 10.1016/j.future.2022.05.014

References

1. Novakovskiy, A., & Yaloveha, I. (2025). Vyznachennia rolii ahentiv velykykh movnykh modelei (LLM) u modeli dyzain-protsesu [Defining the roles of large language model (LLM) agents in the design process model]. *Suchasnyi stan naukovykh doslidzhen ta tekhnolohii v promyslovosti*, (3), 58–72. <https://doi.org/10.30837/2522-9818.2025.3.058>
2. Ministerstvo sotsialnoi polityky Ukrainy. (2020, February 7). Pro zatverdzhennia Metodichnykh rekomendatsii shchodo otsiniuvannia hendernoho vplyvu haluzevykh reform (Nakaz № 86) [On approval of Methodological recommendations for assessing the gender impact of sectoral reforms (Order No. 86)].
3. Wei, J., Wang, X., Schuurmans, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35, 24824–24837.
4. Shojaei, P., Mirzadeh, I., Alizadeh, K., et al. (2025). *The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity*. Apple. <https://doi.org/10.48550/arXiv.2506.06941>
5. Khattab, O., et al. (2024). DSPy: Compiling declarative language model calls into self-improving pipelines. *arXiv*. <https://doi.org/10.48550/arXiv.2310.03714>
6. Yang, C., Wang, X., Lu, Y., et al. (2024). *Large language models as optimizers*. *arXiv*. <https://doi.org/10.48550/arXiv.2309.03409>
7. Hinton, G., Vinyals, O., & Dean, J. (2015). *Distilling the knowledge in a neural network*. *arXiv*. <https://doi.org/10.48550/arXiv.1503.02531>
8. Wu, X., Xiao, L., Sun, Y., Zhang, J., Ma, T., & He, L. (2022). A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems*, 135. <https://doi.org/10.1016/j.future.2022.05.014>

Дата першого надходження статті до видання: 11.02.2026

Дата прийняття статті до друку після рецензування: 17.03.2026

Дата публікації (оприлюднення) статті: 07.05.2026