

DOI: 10.31319/2519-8106.1(54)2026.352417

UDC 004.8

Zhulkovska Inna¹, Candidate of technical sciences, Associate Professor of
Computer Science and Software Engineering Department

Жульковська І.І., кандидат технічних наук, доцент кафедри комп'ютерних наук та
інженерії програмного забезпечення

ORCID: 0000-0002-6462-4299

e-mail: inivzh@gmail.com

Zhulkovskyi Oleg², Candidate of technical sciences, Acting Head of Systems Software Department

Жульковський О.О., кандидат технічних наук, в.о. зав. кафедри програмного забезпечення
систем

ORCID: 0000-0003-0910-1150

e-mail: olalz@ukr.net

Yakovenko Vadym¹, Doctor of technical sciences, Professor of
Computer Science and Software Engineering Department

Яковенко В.О., доктор технічних наук, професор кафедри комп'ютерних наук та
інженерії програмного забезпечення

ORCID: 0000-0002-9582-5990

e-mail: yakovenko@ua.fm

Rudianova Tetiana¹, Candidate of Physical and Mathematical Sciences, Associate Professor of
Computer Science and Software Engineering Department

Рудянова Т.М., кандидат фізико-математичних наук, доцент кафедри комп'ютерних наук та
інженерії програмного забезпечення

ORCID: 0000-0002-8685-4132

e-mail: rudyanova@i.ua

Mala Yuliia¹, Candidate of technical sciences, Associate Professor of
Computer Science and Software Engineering Department

Мала Ю.А., кандидат технічних наук, доцент кафедри комп'ютерних наук та
інженерії програмного забезпечення

ORCID: 0000-0002-2539-4793

e-mail: malaya.ua@gmail.com

Lebedkin Danil², student of of Systems Software Department

Лебьодкін Д.О., здобувач вищої освіти бакалаврського рівня кафедри програмного
забезпечення систем

e-mail: lebedkin041004@gmail.com

¹University of Customs and Finance, Dnipro

Університет митної справи та фінансів, м. Дніпро

²Dniprovsky State Technical University, Kamianske

Дніпровський державний технічний університет, м. Кам'янське

APPROACHES TO ENSURING DATA PRIVACY IN MACHINE LEARNING MODELS

ПІДХОДИ ДО ЗАБЕЗПЕЧЕННЯ КОНФІДЕНЦІЙНОСТІ ДАНИХ У МОДЕЛЯХ МАШИННОГО НАВЧАННЯ

The increasing risks of personal data compromise necessitate formal privacy guarantees in recommendation services. This study investigates the impact of differential privacy on collaborative filtering and matrix factorization, specifically utilizing singular value decomposition. To evaluate robustness against inversion attacks, modifications using additive Laplace and Gaussian noise were implemented. Using RAWG.io metadata, experiments analyzed the non-linear relationship between the privacy budget (ϵ) and accuracy via Precision@10 and RMSE metrics. Results reveal a significant utility-privacy trade-off: while weak protection ($\epsilon = 5.0$) maintains high relevance (Precision = 0.81), the optimal balance is achieved at $\epsilon = 1.0$ – 2.0 (55–65 % accuracy). Maximum privacy ($\epsilon = 0.1$) leads to performance degradation (0.39–0.46) due to noise dominance. SVD-based filtering exhibits higher resilience than content-based approaches by effectively filtering distortions within the latent factor space. The findings confirm the feasibility of differential privacy integration into commercial systems. Future research will focus on adaptive noise mechanisms accounting for data density and matrix sparsity.

Keywords: differential privacy, recommendation systems, ϵ -privacy, Laplace mechanism, collaborative filtering, matrix factorization.

Зростання ризиків компрометації персональних даних у цифрових мережах зумовлює актуальність впровадження формальних гарантій конфіденційності у рекомендаційні сервіси. У роботі проведено комплексне дослідження впливу механізмів диференційної приватності на ефективність алгоритмів колаборативної фільтрації та матричної факторизації, зокрема з використанням методу сингулярного розкладу як базового інструменту побудови латентних представлень. Для оцінки стійкості моделей до атак на інверсію та несанкціоноване відновлення даних реалізовано модифікації методів із використанням адитивних шумів Лапласа та Гауса, інтегрованих у процес навчання.

Експериментальна база дослідження базується на метаданих платформи RAWG.io, що дозволило створити реалістичне середовище моделювання поведінки користувачів. Проведено серію експериментів для аналізу нелінійної залежності між параметром бюджету приватності ϵ та точністю прогнозів за метриками Precision@10 та середньоквадратичної похибки RMSE. Емпірично встановлено наявність суттєвого компромісу між корисністю системи та приватністю: при слабкому рівні захисту ($\epsilon = 5.0$) модель зберігає високу релевантність (Precision = 0.81), тоді як оптимальний баланс досягається в діапазоні $\epsilon = 1.0$ – 2.0 , де точність становить 55–65 %. Граничне посилення конфіденційності ($\epsilon = 0.1$) призводить до деградації показників до 0.39–0.46 через домінування стохастичного шуму.

Доведено вищу стійкість SVD-фільтрації до шумового впливу порівняно з контентно-орієнтованими підходами завдяки здатності алгоритму зберігати структурні закономірності у латентному просторі. Висновки підтверджують можливість інтеграції механізмів диференційної приватності у комерційні рекомендаційні системи за умови коректного калібрування рівня шуму. Перспективи подальших досліджень пов'язані з розробкою адаптивних механізмів додавання шуму з урахуванням цільності та розрідженості даних.

Ключові слова: диференційна приватність, рекомендаційні системи, ϵ -приватність, механізм Лапласа, колаборативна фільтрація, матрична факторизація.

Problem's Formulation

Modern information systems actively integrate ML (Machine Learning), artificial intelligence, and intelligent analytics algorithms into a wide range of application areas—from e-commerce and financial services to healthcare and public administration. The functioning of such systems is based on the processing of large amounts of personal data, including information about user behavior, interests, social interactions, and geolocation [1].

This is particularly evident in recommendation systems that provide personalized interaction with the user. Streaming platforms, online stores, and gaming services analyze viewing history, content ratings, and interaction patterns to generate individual recommendations. Such systems widely use methods of collaborative and content filtering, matrix factorization, and neural network-based models [2]].

At the same time, the data used by ML models often contains confidential information, the leakage or unauthorized use of which poses a significant threat to privacy. In addition to identification attributes, user behavior profiles are processed, including search queries, click history, and consumed content. The combination of such information makes it possible to reconstruct a high-fidelity digital profile of the user, the sensitivity level of which often exceeds that of traditional personal data.

By accumulating vast arrays of behavioral information in e-commerce, social networks, education, and financial services, recommendation systems become potentially vulnerable to privacy attacks, including data reconstruction, inference attacks, and training data leakage.

In this context, a critical scientific problem is the development of approaches and architectures for ML models—specifically recommendation systems—that provide formal mathematical privacy guarantees without significant degradation in predictive accuracy and utility in real-world information environments.

Analysis of recent research and publications

Ensuring privacy in modern ML applications and personalized recommendation systems is based on a combination of mathematically formalized criteria and practical algorithmic approaches aimed at minimizing the risks of sensitive data leakage during model training, deployment, and result sharing. Central to these approaches is DP (Differential Privacy), which guarantees that the addition or removal of a single record from the training set has a limited impact on the model's behavior. This allows for formally bounding the amount of information that can be reconstructed through model access analysis and establishing rigorous mathematical guarantees for user privacy protection [3], [4].

The scientific literature distinguishes between two main approaches to DP implementation: Global and Local DP [5]. Local DP involves adding noise directly on the users' devices, allowing only modified or anonymized data to be transmitted to the server, significantly reducing risks associated with centralized storage. Conversely, Global DP relies on centralized data collection followed by the application of noise injection and gradient clipping algorithms; this provides higher model accuracy, but requires trust in the central server.

Distributed architectures, particularly Federated Learning [6], occupy an intermediate position, combining local model training on user devices with centralized aggregation. In such systems, controlled noise is added during the update exchange phase, keeping raw data local while maintaining the consistency of the global model.

In the process of training models with DP, gradients play a key role, as they determine the direction and magnitude of parameter updates to minimize the loss function. To prevent sensitive information leakage, controlled noise is added to the gradients, and a norm clipping mechanism is applied. The level of protection is determined by the privacy budget (ϵ , δ), which formalizes the permissible level of potential information disclosure regarding individual records. Lower ϵ values ensure higher privacy but may lead to reduced predictive accuracy, whereas increasing ϵ improves recommendation quality at the cost of weakened confidentiality guarantees.

These principles form the basis of the DPML (Differential Privacy in Machine Learning) approach, which integrates protection mechanisms directly into the model training process. In recommendation systems, this is specifically realized by modifying matrix factorization algorithms, where the perturbation of parameters and updates formally bounds the probability of identifying individual users in the final model.

Studies into the practical aspects of DPML confirm the effectiveness of integrating stochastic mechanisms, such as Laplace and Gaussian distributions, directly within the model optimization process [4], [8]. In recommendation systems, the most common approaches combine DP-SGD (Differentially Private Stochastic Gradient Descent) with matrix factorization methods to protect both gradients and model parameters.

At the same time, given the parallel evolution of de-anonymization methods and inference attacks, the development of attack-resistant architectures for recommendation systems remains an open question. Maintaining the balance between formal privacy guarantees and high recommendation quality continues to be one of the key scientific and applied challenges in this field [9]—[12].

Formulation of the study purpose

Modern personalized recommendation systems face a fundamental tension between the need to protect sensitive user data and the requirement to maintain high recommendation relevance. Addressing this conflict requires a thorough investigation of privacy-preserving methods and mechanisms, as well as their impact on predictive quality.

The objectives of this study are:

1. To evaluate the performance and robustness of privacy-preserving methods in ML models and recommendation systems.
2. To analyze the trade-off between the level of user data privacy guarantees and prediction accuracy.
3. To develop practical guidelines for integrating privacy-preserving approaches into real-world personalized recommendation services.

Presenting main material

The theoretical and mathematical foundations of modern recommendation systems are based on the formalization of interactions between a set of users and a set of items, where the interaction matrix serves as the central element. Formally, the interaction matrix is defined as $R \in \mathbb{R}^{|U| \times |I|}$, where R_{ui} represents the degree of interaction between user u and item i . Most real-world recommendation systems, however, operate under conditions of extreme data sparsity, which complicates the direct application of classical statistical methods and necessitates approaches capable of generalizing from limited observations.

Given these challenges, the following conceptual architecture illustrates how privacy-preserving mechanisms are integrated into the recommendation process (Fig. 1). In this architecture, raw user interaction data are first processed through a DP layer, where controlled noise is injected according to the privacy budget parameter ϵ . The resulting perturbed data are then projected into a latent space using SVD-based (Singular Value Decomposition) matrix factorization, enabling the extraction of stable behavioral patterns. These latent representations are subsequently used for personalized recommendation generation. This architectural scheme emphasizes the crucial role of latent factors in preserving the predictive accuracy of the system while mitigating the potential negative effects of privacy-preserving perturbations.

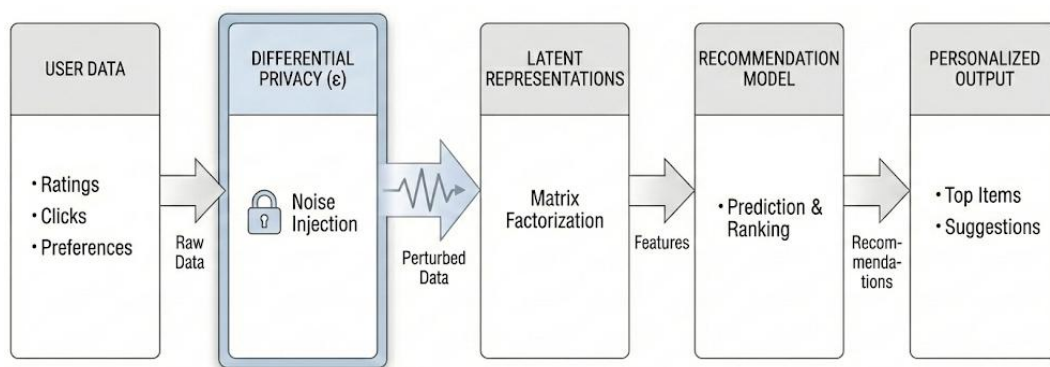


Fig. 1. Conceptual architecture of a personalized recommendation system with integrated DP mechanisms

The practical complexity of this problem stems from the fact that most real-world recommendation systems operate under conditions of extreme data sparsity. The number of available items significantly exceeds the volume of interactions for any individual user; consequently, the vast majority

of the matrix elements remain undefined. This property significantly complicates the application of classical statistical methods and necessitates the use of models capable of generalizing information based on a limited number of observations.

A key approach to solving this problem is the use of latent factors—hidden parameters that reflect fundamental patterns in user choices and item properties. Latent factors are not directly observed but are derived from accumulated behavioral data. They allow for the representation of both users and items in a shared multidimensional feature space, where the distance or angle between vectors reflects the degree of preference alignment.

In this case, each user and item correspond to vectors $\mathbf{p}_u, \mathbf{q}_i \in \mathbb{R}^k$, and the predicted interaction is approximately defined by the dot product $\hat{R}_{ui} = \mathbf{p}_u^T \mathbf{q}_i$.

From a mathematical perspective, this approach provides dimensionality reduction, replacing the analysis of a large sparse matrix with a compact representation in the form of fixed-length vectors. This not only reduces the computational complexity of algorithms but also enhances the model's generalization capability, allowing for recommendations even in the absence of direct matches in the rating history.

At the same time, it should be noted that latent factors accumulate the most informative part of a user's personal profile. They reflect stable behavioral patterns rather than individual actions, which makes them particularly sensitive from a privacy standpoint. Consequently, in modern recommendation systems, protecting latent representations is considered a critical requirement for ensuring privacy and resilience against inference attacks.

In collaborative recommendation models, evaluating the similarity between users or items plays a fundamental role. The most common metric is cosine similarity, defined as the normalized dot product of the respective interaction vectors [2], [13].

Cosine similarity is defined as:

$$\text{sim}(u, v) = \frac{\mathbf{r}_u \cdot \mathbf{r}_v}{\|\mathbf{r}_u\| \|\mathbf{r}_v\|}.$$

This approach allows for evaluating the direction of preferences rather than the absolute magnitude of activity, which is fundamentally important in systems with an uneven distribution of ratings.

The calculated similarity values are used to form a set of nearest neighbors, which serves as the basis for predicting missing ratings. The prediction is determined as a weighted average of relevant users' ratings, where the weights are proportional to the degree of profile similarity. This scheme enables the efficient integration of local information about user preferences into the global recommendation process.

In practical implementations of collaborative algorithms, a significant aspect is the handling of anomalous behavior. Users with atypical or extreme rating patterns can substantially influence prediction results. To address this, multi-level outlier filtering and dynamic adjustment of connection weights between profiles are applied. This results in the formation of a dynamic similarity matrix capable of responding promptly to shifts in user interests and global system trends.

Unlike collaborative methods, content-based approaches are founded on the analysis of the intrinsic characteristics of items. One of the fundamental tools for such analysis is the TF-IDF (Term Frequency — Inverse Document Frequency) model, which enables the formalization of textual item descriptions into numerical vectors. The algorithm accounts for both the local significance of a term within a specific description and its global prevalence across the entire dataset, ensuring the extraction of semantically informative features.

The application of TF-IDF in conjunction with matrix factorization enables the design of a hybrid recommendation model capable of efficiently processing both interaction history and descriptive content attributes [14], [15]. Within the scope of this study, the SVD algorithm was selected as the mathematical foundation for matrix factorization. Using SVD as a specific yet most representative case of factorization enables the decomposition of a sparse rating matrix into the product of orthogonal latent factor matrices. This not only provides a compact data representation but also allows for the extraction of the most significant singular values, which correspond to stable user behavior patterns.

Such an approach is particularly effective for addressing the cold start problem, where new items lack a sufficient number of ratings. Furthermore, SVD factorization creates favorable conditions for integrating DP mechanisms. Manipulating latent factors allows for the addition of controlled stochastic noise without directly affecting interpretable user actions.

To ensure formal privacy guarantees during model training, DP mechanisms are embedded directly into the latent factor learning process. Fig. 2 illustrates the conceptual architecture of a privacy-preserving recommendation system, where controlled stochastic noise is injected into latent representations, enabling protection of individual user contributions while preserving the overall predictive structure of the model.




Fig. 2. Integration of DP mechanisms into a matrix-factorization-based recommendation system

Consequently, fundamental interaction patterns are preserved even after perturbation, while the risk of reconstructing individual data is substantially reduced. This occurs because the noise is distributed within a reduced-dimensionality space, neutralizing the possibility of identifying specific user transactions.

To provide formal data protection guarantees, this study employs the DP principle. According to this framework, an algorithm is considered ϵ -differentially private if a change in a single record in the input dataset does not lead to a significant increase in the probability of obtaining any specific output. The parameter ϵ acts as a quantitative measure of the privacy budget and defines the trade-off between the level of protection and the accuracy of the results.

Formally, this requires satisfying the following inequality:

$$\Pr[M(D) \in S] \leq e^\epsilon \Pr[M(D') \in S],$$

where D and D' differ by a single record.

The practical implementation of DP involves adding stochastic noise, specifically through the Laplace or Gaussian mechanisms. The noise intensity is determined by the global sensitivity of the target function, which reflects the maximum impact of a single user record on the algorithm's output. DP mechanisms can be applied at various stages—ranging from input data perturbation to gradient modification during model training.

Experimental results indicate that various components of a hybrid recommendation system exhibit different degrees of sensitivity to noise. While content-based methods demonstrate higher stability, collaborative algorithms prove to be significantly dependent on the ϵ parameter due to their direct operation with interaction matrices.

Integrated analysis of Fig. 3 reveals a pronounced privacy-utility trade-off between recommendation accuracy and the model's prediction error. Under strict privacy constraints ($\epsilon = 0.1$), the system exhibits the largest utility degradation: Precision@10 reaches its minimum value (0.46), while RMSE peaks at 1.00, which is indicative of intensive noise injection and substantial signal distortion. As the privacy budget increases, both metrics improve monotonically, and a balanced operational regime emerges within the range $\epsilon = 1.0$ – 2.0 . In this interval, the model maintains sufficiently high recommendation relevance (Precision@10 > 0.65) while simultaneously achieving a notable reduction in prediction error (RMSE < 0.8), representing the most favorable compromise between data protection and predictive quality. Further relaxation of privacy constraints ($\epsilon \rightarrow 5.0$) leads to metric stabilization rather than proportional gains, with precision saturating at 0.81 and RMSE decreasing marginally to 0.56, which indicates diminishing returns. This behavior confirms that the use of latent representations obtained via SVD factorization effectively attenuates stochastic perturbations, preserving robust model functionality even under substantial noise-induced distortion.




Fig. 3. Dependence of recommendation accuracy and prediction error on the DP budget ϵ

The evidence suggests that latent factor models inherently mitigate the impact of data perturbation, ensuring that privacy-preserving measures do not degrade the overall recommendation logic. These findings prove the feasibility of integrating DP mechanisms into recommendation systems [16], providing mathematically sound guarantees for personal data protection without compromising the service's consumer value.

Conclusions

The study examined the methodological and applied aspects of ensuring confidentiality within the framework of modern personalized recommendation systems. Based on the analysis of collaborative filtering using SVD factorization and the content-based approach, the key patterns of DP impact on prediction quality across various operational scenarios were identified and mathematically justified.

Experimental evaluation, conducted using real-world metadata from the RAWG.io platform, demonstrated that the implementation of DP mechanisms achieves a high level of user data protection without a critical loss in the relevance of personalized results. Specifically, the Precision@10 analysis confirmed the robust performance of the system within the ϵ parameter range of 1.0—5.0, where the model effectively filters calibrated stochastic noise in latent spaces. The analysis of noise dispersion further verified that these perturbations remain unbiased, preserving the overall data topology. This is attributed to the fact that latent factor vectors accumulate the most stable behavioral patterns, which prove to be more resistant to Laplace perturbation compared to raw ratings. The experiments identified $\epsilon = 1.0$ as the optimal privacy budget for balancing rigorous data masking with recommendation utility.

The comparative analysis revealed a significant differentiation in recommendation quality depending on the chosen algorithmic basis and the specifics of the input data. It was established that personalized collaborative predictions exhibit substantially higher robustness to external noise compared to content-based methods for finding similar items. The lower effectiveness of the latter within the studied dataset is due to the high degree of sparsity in textual attributes and the complexity of formalizing the semantic similarity of game projects based solely on genres and metadata. This confirms that the analysis of direct user interactions serves as a more reliable signal for building accurate predictions, even under strict privacy budget constraints.

The experiments proved that the concept of DP is a viable and practically applicable solution that harmonizes legal requirements for personal data protection with business needs for effective content personalization.

Future research will focus on developing adaptive noise injection mechanisms capable of accounting for individual data sparsity across different system segments. Another promising direction is improving content-based algorithms by integrating more sophisticated neural network models for natural language processing to enhance semantic analysis quality. Furthermore, it is necessary to investigate dynamic privacy budget allocation in high-load real-time systems, where maintaining model relevance requires continuous data updates.

The implementation of these approaches will achieve an optimal balance between mathematically guaranteed security and high consumer value of personalized services in future developments, creating a solid foundation for the next generation of privacy-preserving intelligent systems.

References

- [1] Shokri R., Stronati M., Song C., Shmatikov V. (2017). Membership inference attacks against machine learning models. In: Proc. 2017 IEEE Symposium on Security and Privacy (SP). San Jose, CA, USA. P. 3–18 [in English]. DOI: <https://doi.org/10.1109/SP.2017.41>
- [2] He X., Liao L., Zhang H., Nie L., Hu X., Chua T.-S. (2017). Neural collaborative filtering. In: Proceedings of the 26th International Conference on World Wide Web (WWW '17). Perth, Australia. P. 173–182 [in English]. DOI: <https://doi.org/10.1145/3038912.3052569>
- [3] Dwork C., Roth A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*. Vol. 9, No. 3–4. P. 211–407 [in English]. DOI: <https://doi.org/10.1561/04000000042>
- [4] Abadi M., Chu A., Goodfellow I. et al. (2016). Deep learning with differential privacy. In: Proc. 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16). New York, NY, USA. P. 308–318 [in English]. DOI: <https://doi.org/10.1145/2976749.2978318>
- [5] Wang T., Fan L., Jin H., Wang P., Yang Q. (2020). A comprehensive survey on local differential privacy: Theory, methods, and applications. *Sensors*. Vol. 20, No. 12. Art. 3436 [in English]. DOI: <https://doi.org/10.3390/s20123436>
- [6] Wu C., Wu F., Lyu L. et al. (2022). A federated graph neural network framework for privacy-preserving personalization. *Nat Commun*. Vol. 13. Article 3091. DOI: <https://doi.org/10.1038/s41467-022-30714-9>
- [7] Wang J., Guo S., Xie X., Qi H. (2022). Protect privacy from gradient leakage attack in federated learning. In: Proceedings of the 41st IEEE International Conference on Computer Communications (INFOCOM '22). P. 580–589 [in English]. DOI: <https://doi.org/10.1109/INFOCOM48880.2022.9796841>
- [8] Müllner P. (2023). Differential privacy in collaborative filtering recommender systems. *Frontiers in Big Data*. Vol. 6. Art. 1151053 [in English]. DOI: <https://doi.org/10.3389/fdata.2023.1249997>
- [9] Milano S., Taddeo M., Floridi L. (2020). Recommender systems and their ethical challenges. *AI & Society*. Vol. 35. P. 957–967 [in English]. DOI: <https://doi.org/10.1007/s00146-020-00950-y>
- [10] Yang Z., Zhang X., He X. et al. (2021). Membership inference attacks against recommender systems. In: Proc. 2021 ACM Conference on Computer and Communications Security (CCS). Seoul, South Korea. P. 864–879 [in English]. DOI: <https://doi.org/10.1145/3460120.3485365>
- [11] Zhu L., Liu Z., Han S. (2019). Deep leakage from gradients. In: Advances in Neural Information Processing Systems (NeurIPS '19). P. 1–11 [in English]. DOI: <https://doi.org/10.48550/arXiv.1906.08935>
- [12] Fredrikson M., Jha S., Ristenpart T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15). Denver, CO, USA. P. 1322–1333 [in English]. DOI: <https://doi.org/10.1145/2810103.2813677>
- [13] Hitaj B., Ateniese G., Perez-Cruz F. (2017). Deep models under the GAN: Information leakage from collaborative deep learning. Proc. 2017 ACM SIGSAC CCS. Dallas, TX, USA. P. 603–618. DOI: <https://doi.org/10.1145/3133956.3134012>
- [14] Gomez-Uribe C.A., Hunt N. (2016). The Netflix recommender system: algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*. Vol. 6, No. 4. Art. 13 [in English]. DOI: <https://doi.org/10.1145/2843948>
- [15] Zhang S., Yao L., Sun A., Tay Y. (2019). Deep learning based recommender system: a survey and new perspectives. *ACM Computing Surveys*. Vol. 52, No. 1. Art. 5 [in English]. DOI: <https://doi.org/10.1145/3285029>
- [16] Zhulkovskii O., Panteikov S., Zhulkovskaya I. (2022). Information-modeling forecasting system for thermal mode of top converter lance. *Steel in Translation*. Vol. 52, No. 5, P. 495–502. [in English]. DOI: <https://doi.org/10.3103/S0967091222050138>

Список використаної літератури

1. Shokri R., Stronati M., Song C., Shmatikov V. Membership inference attacks against machine learning models. Proc. 2017 IEEE Symposium on Security and Privacy (SP). San Jose, CA, USA. 2017. P. 3–18. DOI: <https://doi.org/10.1109/SP.2017.41>
2. He X., Liao L., Zhang H., Nie L., Hu X., Chua T.-S. Neural collaborative filtering. Proceedings of the 26th International Conference on World Wide Web (WWW '17). Perth, Australia. 2017. P. 173–182. DOI: <https://doi.org/10.1145/3038912.3052569>
3. Dwork C., Roth A. The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science. 2014. Vol. 9, № 3–4. P. 211–407. DOI: <https://doi.org/10.1561/04000000042>
4. Abadi M., Chu A., Goodfellow I. et al. Deep learning with differential privacy. Proc. 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16). New York, NY, USA. 2016. P. 308–318. DOI: <https://doi.org/10.1145/2976749.2978318>
5. Wang T., Fan L., Jin H., Wang P., Yang Q. A comprehensive survey on local differential privacy: Theory, methods, and applications. Sensors. 2020. Vol. 20, № 12. Article 3436. DOI: <https://doi.org/10.3390/s20123436>
6. Wu C., Wu F., Lyu L. et al. A federated graph neural network framework for privacy-preserving personalization. Nat Commun. 2022. Vol. 13. Article 3091. DOI: <https://doi.org/10.1038/s41467-022-30714-9>
7. Wang J., Guo S., Xie X., Qi H. Protect privacy from gradient leakage attack in federated learning. Proceedings of the 41st IEEE International Conference on Computer Communications (INFOCOM '22). 2022. P. 580–589. DOI: <https://doi.org/10.1109/INFOCOM48880.2022.9796841>
8. Müllner P. Differential privacy in collaborative filtering recommender systems. Frontiers in Big Data. 2023. Vol. 6. Article 1151053. DOI: <https://doi.org/10.3389/fdata.2023.1249997>
9. Milano S., Taddeo M., Floridi L. Recommender systems and their ethical challenges. AI & Society. 2020. Vol. 35. P. 957–967. DOI: <https://doi.org/10.1007/s00146-020-00950-y>
10. Yang Z., Zhang X., He X. et al. Membership inference attacks against recommender systems. Proc. 2021 ACM Conference on Computer and Communications Security (CCS). Seoul, South Korea. 2021. P. 864–879. DOI: <https://doi.org/10.1145/3460120.3485365>
11. Zhu L., Liu Z., Han S. Deep leakage from gradients. Advances in Neural Information Processing Systems (NeurIPS '19). 2019. P. 1–11. DOI: <https://doi.org/10.48550/arXiv.1906.08935>
12. Fredrikson M., Jha S., Ristenpart T. Model inversion attacks that exploit confidence information and basic countermeasures. Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15). Denver, CO, USA. 2015. P. 1322–1333. DOI: <https://doi.org/10.1145/2810103.2813677>
13. Hitaj B., Ateniese G., Perez-Cruz F. Deep models under the GAN: Information leakage from collaborative deep learning. Proc. 2017 ACM SIGSAC CCS. Dallas, TX, USA. 2017. P. 603–618. DOI: <https://doi.org/10.1145/3133956.3134012>
14. Gomez-Urbe C. A., Hunt N. The Netflix recommender system: algorithms, business value, and innovation. ACM Transactions on Management Information Systems. 2016. Vol. 6, № 4. Article 13. DOI: <https://doi.org/10.1145/2843948>
15. Zhang S., Yao L., Sun A., Tay Y. Deep learning based recommender system: a survey and new perspectives. ACM Computing Surveys. 2019. Vol. 52, № 1. Article 5. DOI: <https://doi.org/10.1145/3285029>
16. Zhulkovskii O., Panteikov S., Zhulkovskaya I. Information-modeling forecasting system for thermal mode of top converter lance. Steel in Translation. 2022. Vol. 52, No. 5, P. 495–502. DOI: <https://doi.org/10.3103/S0967091222050138>

Надійшла до редколегії 06.01.2026

Прийнята після рецензування 14.01.2026

Опублікована 22.01.2026