Scientific Visualization, 2024, volume 16, number 5, pages 109 - 119, DOI: 10.26583/sv.16.5.08

The Impact of Input Data Density on the Performance of Graphic Neural Networks

N.A. Bondareva[1,A]

Keldysh Institute of Applied Mathematics RAS

[1] ORCID: 0000-0002-7586-903X, nicibond9991@gmail.com

Abstract

The paper provides a brief overview of generative neural networks and considers the role of information in training generative neural networks. In the digital environment, each object is surrounded by a vast information field, including unordered information and a set of references to it. The density of the object's information field determines the ability of technologies such as artificial intelligence to recreate its image based on the collected data. The more data is available, the more accurately and completely the digital image can be recreated. The paper considers a number of problems arising from the use of text-to-image networks and possible methods for solving them. The article considers various aspects of the role of personal data and possible ethical and social consequences in the era of generative technologies, as well as the prospects and risks of further development of generative neural networks in specialized areas such as medicine and manufacturing. The rapid development of neural network technologies can have a significant impact on education and social phenomena.

Keywords: Machine learning, computer vision and pattern recognition, neural network, computer graphics, information field density, Text-to-image.

1. Introduction

Currently, the field of neural network technologies is developing rapidly, acquiring more and more skills and capabilities every day. Particularly popular are generative diffusion models, which are the basis of most text neural networks that can collect, analyze and generate text information on request, and graphic ones that can process media content in various ways, from animating photos to automatically creating images and videos on a text request. The results of their work are used in many industries, from media to medicine, but rapid progress also causes social changes. Previously, generative adversarial neural networks (GAN) were considered a promising alternative to diffusion models, but they turned out to be less effective for generating images from text and unstable in training [1].

Diffusion models are iterative algorithms that transform random noise into an image. An example is the DDPM (Denoising Diffusion Probabilistic Model) [2], trained on thousands of images to which noise is successively added. The model learns to remove this noise, improving the quality of the image. If a trained model is applied to random noise, it can create a new image, gradually clearing it of noise. The figure shows an example where the user specifies a schematic drawing. The image is noisy, and then the model reconstructs it with high accuracy through a process of back diffusion.

Graphic neural networks interpret the linguistic structure of a query, process it, and generate realistic visual images. They manage multiple objects, their attributes, and spatial relationships, establishing the correct connections between object characteristics. The basis for such tasks are diffusion neural network models [2,3], which appeared in 2015 but gained popularity after the work of [2]. Today, they achieve impressive results in generating and modifying images, such as generating images, music, and video from a text query (text-to-image), restoring details (inpainting), removing objects, and increasing resolution (super-resolution).

Text-to-image models use a linguistic construct (a textual query) to guide their processing. Language models trained on pairs of images and texts understand the content of both types of data. For example, the CLIP model (Contrastive Language — Image Pre-training) from OpenAI transforms images and texts into a common latent vector space, where vectors represent a set of values. In such a space, you can find the closest images to a text query simply by manipulating vectors. The Latent Model Diffusion [4], introduced in 2021, trains image generation from directional noise using a latent space for texts and images. The same principles are applied in Stable models Diffusion, Imagen and other large neural networks for converting text into images. The basic principles of their operation are described in [5-8].

2. The Role of Information in Training Generative Neural Networks

Generative neural networks such as Generative Adversarial Networks (GAN), Variational Autoencoders (VAEs) and transformers (e.g. GPT) are complex systems that can create new data that is indistinguishable from real data. The role of information in these models is multifaceted and critical to their successful operation and training. Based on the principles of machine learning, data becomes the foundation on which the training model is built, allowing it to generalize, interpolate, and extrapolate patterns.

A separate task in addition to training the neural network is the formation of training datasets [9]. The dataset is the main source of information on which the neural network builds its model. The model learns from the dataset, analyzing it, identifying patterns and regularities that it subsequently uses for generation. Its quality and diversity directly affect the capabilities of the generated model: the better the data is prepared, the less time it will take to debug the model, train it, find and eliminate recognition inaccuracies.

Among the main criteria of quality The following can be distinguished from the dataset:

1. Completeness of data. This means that the data sets are sufficient in size, depth, and breadth. The data set should contain enough parameters or features so that there are no edge cases left uncovered. Incompleteness results in either the impossibility of analysis or the need to rely on some assumptions or presumptions regarding the missing information.
2. Accuracy. The data should be as close as possible to the real conditions in which the neural network model will operate.
3. Correctness and correctness. This point implies the correspondence of the data to reality and the correctness of their interpretation, as well as the correspondence of the format and annotations of the data in the dataset with those in which the framework and architecture of the neural network model operate.
4. Uniformity - the values of all attributes should be comparable across all data. Unevenness or the presence of outliers in data sets negatively affect the quality of training data.
5. Having separate datasets for training, validation and testing.

Large language models (LLMs) have revolutionized natural language processing in tasks such as reading comprehension, reasoning, and language generation. They are powerful tools in the field of natural language processing (NLP). Such models are trained on huge text datasets and are able to capture complex patterns and nuances of language.

The basic principle of LLM is to use the transformer architecture proposed in [10]. Transformers allow us to identify and analyze dependencies between words in a sentence, regardless of their position. This improves the quality of text generation and context understanding.

LLMs are widely used in many types of generative neural networks, including text - to - image , text - to - video , and text - to - 3D , as well as in a variety of text and text query processing, from automatic translation to generating fully coherent text or code. Models are trained on many thousands of gigabytes of text data, including books, articles, websites, and other various text resources. The amount of information contained in an LLM model can be characterized as extremely high. For example, OpenAI 's GPT-3 model has 175 billion parameters, while its predecessor GPT-2 had only 1.5 billion parameters.

Kandinsky GPU neural network model was trained on the SberCloud ML Space platform for two months, spending 20,352 GPU-V100 days. It was trained on a dataset of 60 million pairs of images and text descriptions, which was subsequently reduced to 28 million. Such well-known datasets as ConceptualCaptions [11] (a dataset containing over 3 million images with natural language captions, the raw descriptions for which are collected from the Internet) and YFCC100m [12] (the largest publicly available and freely used multimedia collection, containing metadata for about 99.2 million photographs), translated into Russian, were used in the training. The first stage of training consisted of 250 thousand iterations.

However, the capabilities of neural networks are not limited to natural language processing and image generation. For specialized tasks, such as manufacturing or medicine, there is a need for special neural networks trained on professionally oriented datasets. Such neural networks must be able to perceive narrow-profile jargon and scientific and technical terms. This is necessary to prevent possible ambiguous interpretations. Given the huge amount of accumulated data and the availability of specialized archives for many areas, the creation of graphic neural networks with a specific focus is only a matter of time. Their potential application opens up broad opportunities for analyzing and comparing various types of data, as well as for their visualization in an accessible and clear form.

Such neural networks can also be useful in teaching aids. For example, a neural network can reflect the typical condition of an organ or tissue in the presence of a certain set of symptoms mentioned in the request. If the description text indicates any pathology, visualization can help highlight its features, which helps make the right diagnostic decision.

In the manufacturing field, huge databases of technical drawings and standard formalized and detailed 3D models provide a potential opportunity to train specialized neural networks oriented to strict formalization of the query and limited subject matter. Such specialized neural networks can be used to generate CAD files using text hints, as shown in Figure 1, to reconstruct a 3D model from drawings, and to evaluate and compare 3D model trees to identify identical models despite differences in their construction. The generated models can be imported into the CAD program of choice, or specialized Text - to -CAD generators can be created without building and maintaining an infrastructure.

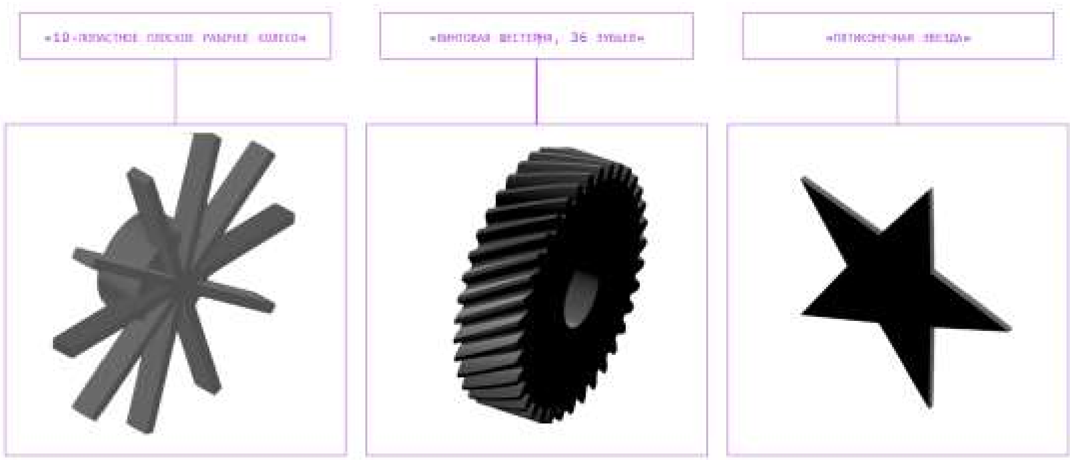

Figure 1 - Text-to-CAD, an interface for creating CAD files using text prompts [13]

In addition, modern text-to-image neural networks, which allow creating images or 3D models based on a text request, can be used in the production process to obtain a preliminary visual appearance of a part, which can then be reworked in accordance with the designer's vision. An example of such a neural network is shown in Figure 2. A preliminary concept, which does not require significant costs and is provided in an unlimited number of options, can significantly reduce the labor intensity and costs of creating prototypes in the process of research and development (R&D).

Figure 2 - Example of creating a CAD system using text prompts

3. Information field of the object

Each object in the digital environment has a so-called information field. The information field of an object is defined as the entire volume of unordered information associated with the sought object and the totality of references in the digital environment. In other words, this is the amount of open and public information that surrounds the sought object and allows its image to be recreated artificially.

The information field includes all references to the object in the digital environment. These may be:

1. Texts: Articles, reviews, comments, posts on social networks, blogs, scientific papers.
2. Multimedia: Photos, videos, graphics, audio recordings.
3. Structured information: Databases, tables, questionnaires, surveys.
4. Metadata: creation time, authorship, location, and other characteristics that help in identifying and processing information about an object.
5. Contextual relationships: relationships between an object and other objects, events, or influencing factors.

Density characterizes the ability of a technology, such as artificial intelligence, to recreate an image of an object based on the data collected. The more data is available, the more accurately and completely the digital image of the object can be recreated.

The density of the information field should directly correlate with such factors as, for example, the frequency of mentions of an object in various media, the diversity of information sources, the depth and detail of the data provided.

Thus, objects that are most frequently and diversely mentioned in public sources will have a high density of the information field, since they are often in the center of attention of the media and the public. But objects or people that are less well-known and less frequently mentioned in public sources will have a lower density. The simplest visualization model that can generally reflect such a representation is the tag cloud shown in Figure 3. It reflects the information distribution of words in one of the sections of this article: the more often a word was mentioned in the text, the larger it is in size.

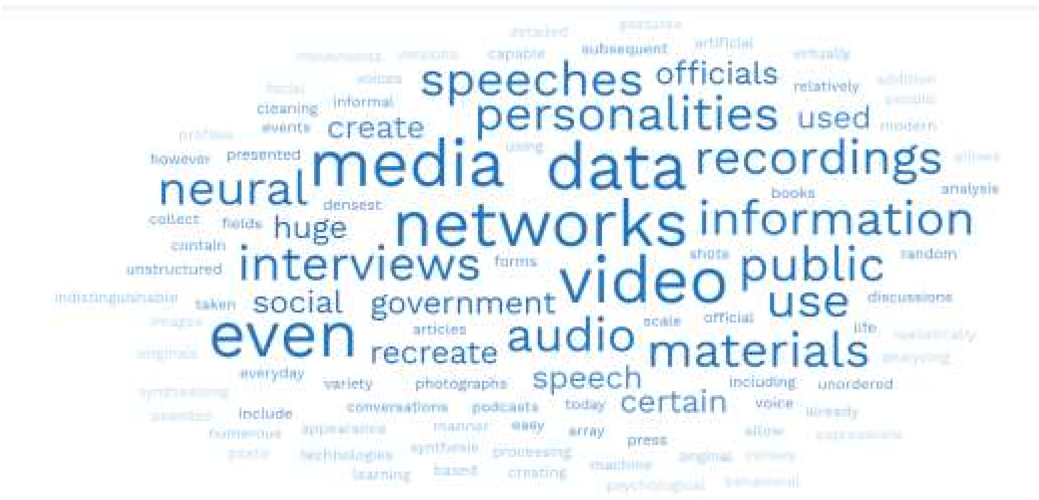

Figure 3 - An example of visualization of the density of the information field in its simplest embodiment (a tag cloud reflecting the most frequently repeated terms in the article).

The densest information fields today are those of media personalities and government officials. Information about them is presented on a huge scale and in a variety of forms, including video materials, voice recordings, photographs, books, articles in the press, discussions on social networks, and much more. Video materials include both official speeches and interviews, as well as random shots taken at public events or even in everyday life. Audio recordings can include speeches, interviews, podcasts, and even informal conversations.

This is a huge array of unorganized and unstructured information, which is nevertheless relatively easy to collect. Subsequent analysis, data cleaning and processing allow this data to be used to recreate artificial appearance, speech patterns, type of eo and audio materials.

Modern speech synthesis technologies are already capable of creating voices of media personalities that are virtually indistinguishable from the original, based on numerous audio recordings. Using machine learning methods and neural networks, it is also possible to recreate a video version of these people, synthesizing images that will convey the facial expressions, gestures, and movements of the originals with maximum realism.

In addition, all this data allows us to create detailed psychological and behavioral profiles of media personalities and government officials. By analyzing their public speeches, interviews, and social media posts, we can identify their preferences, beliefs, and motivations. For example, even the use of certain words and phrases can provide insight into a person's communication style, emotional state, and even professional competence. This data can be used not only to create accurate simulations, but also to predict the behavior of these individuals in certain situations.

The use of such potential capabilities of neural networks can be different. The article [14] considers examples of known cases of malicious use of neural networks, from fraud to manipulation of public opinion.

Currently, in most popular neural networks, especially those that provide their services on a paid basis, developers are taking a wide range of special measures aimed at reducing potential harmful influence and ensuring public safety. Developers are introducing more and more restrictions regarding the use of images of famous government and media figures in images generated by neural networks. However, since there remains the possibility of local user training of individual neural network models, it can be assumed that the problem will remain relevant for a long time.

At the same time, the trend of creating digital doubles of real people, both living and long dead, is gaining strength and popularity. The idea of a "digital pantheon" is not new, but now, with the ability to train a neural network model on footage from newsreels, excerpts from personal correspondence, and collected works of famous historical figures, there is a risk of new methods of manipulating public opinion and polluting the information space and educational system with false or artificially generated information.

This phenomenon can be called the creation of a so-called digital pseudo-personality, which can reproduce the speech and way of thinking of a certain person with a certain degree of reliability. Currently, there is a service for creating such a pseudo-personality for public figures or those doing business [15]. To do this, you will need to upload your voice, digital samples of your appearance (photos and videos), samples of personal and business correspondence, as well as examples of texts in various styles to the database. It is claimed that such a digital pseudo-personality will be able to imitate the communication style of its original and negotiate on its behalf (for example, with clients).

There are a number of risks with this idea:

1) Strict confidentiality of personal data is necessary. If the developer allows them to "float away" into the free Internet, then the user will no longer be able to control the further development of their potential doubles.

2) Fake facts and statements. The user may encounter the fact that unknown facts appear in his biography, which are in fact the product of the digital model generation. He may also be credited with words that he did not say.

3) Manipulation and fraud. The collection and storage of biometric and behavioral data in private commercial organizations carries the risk that the data will end up in the hands of malicious parties. If a digital copy is used by someone other than the original, it is easy to spread misinformation or impersonate someone else.

4) Legal liability. The lack of clarity in the legal regulation of the creation and use of digital copies can create legal vacuums that will be used for illegal purposes, since it is unclear who will be responsible for the words and actions of a digital copy of a real person.

5) Ethical aspects. In August 2024, a video from a supposedly "dead" person was first distributed, who continued to exist in digital form and maintain a blog. The video was recognized as a fake, and the original person turned out to be alive and declared the event an art performance [16], however, in society this created an information precedent for using the face and personal data of a deceased person to reproduce his digital pseudo-personality and further manipulations on his behalf. Abuse of such actions will lead to violations of the rights to privacy and personal data.

When using the personality of famous historical figures, who have a dense information field that allows the creation of a digital pseudo-personality, the following series of problems may arise:

1) Distortion of historical truth, inaccuracies and falsifications. It will become difficult to separate generated statements (especially if they become catchphrases) from real ones, and it will also be difficult to establish the authenticity of a statement, which will increase the risk of manipulating public opinion in historical and political disputes. The use of such copies to interpret historical events can significantly change the perception and understanding of history, which will not always correspond to reality.

2) Political and social manipulation. Digital copies of historical figures can be used for propaganda and political manipulation, and incorrect representation of historical figures (especially if they were controversial figures in their historical era) can cause social tensions.

3) Educational risks. Future generations risk encountering the phenomenon of false authenticity, when a digital copy is perceived as a reliable representation of the personality of a historical figure, which in turn will lead to superficial perception and mass distortions of real facts and the cultural context of a particular era.

Meanwhile, as the internet becomes increasingly populated with neural network- generated data, new problems and potential risks arise from the consequences of such training. Training on data generated by the AI itself can lead to errors being accumulated and redundant repetitions. This comes with the risk of introducing artifacts that are difficult to notice and correct. The paper [17] examines how using model-generated content in training causes irreversible defects in the resulting models, where the tails of the original content distribution disappear. The authors call this effect model collapse and demonstrate that it can occur in variational autoencoders , Gaussian mixture models, and LLMs.

Experts around the world have different opinions [18] on the future of neural networks and the steps that can be taken to minimize future collapse, but many agree that the purity and relevance of data will be key in training future generations of neural networks.

4. The Role of Personal Data in the Age of Generative Neural Networks

In the era of the flourishing of generative neural networks, the personal information about a person that he most often discloses about himself acquires a special role. Local models can be trained even on the small information field that consists of social networks and other sources.

Such data may include:

- photos on social networks;
- video;
- audio recording of voice - from video or telephone conversations
- personal posts and blogs
- diaries
- signature samples

Nowadays, people are constantly being filmed regardless of their own desire - from street video cameras to accidentally getting into other people's videos posted on the Internet. Collecting such information in a personalized manner is not easy, it requires a lot of resources, but, nevertheless, such materials can become training material for new neural network models. Figure 4 shows an annotated photograph from the Diversity training dataset in Faces from IBM, prepared for training the process of facial recognition.

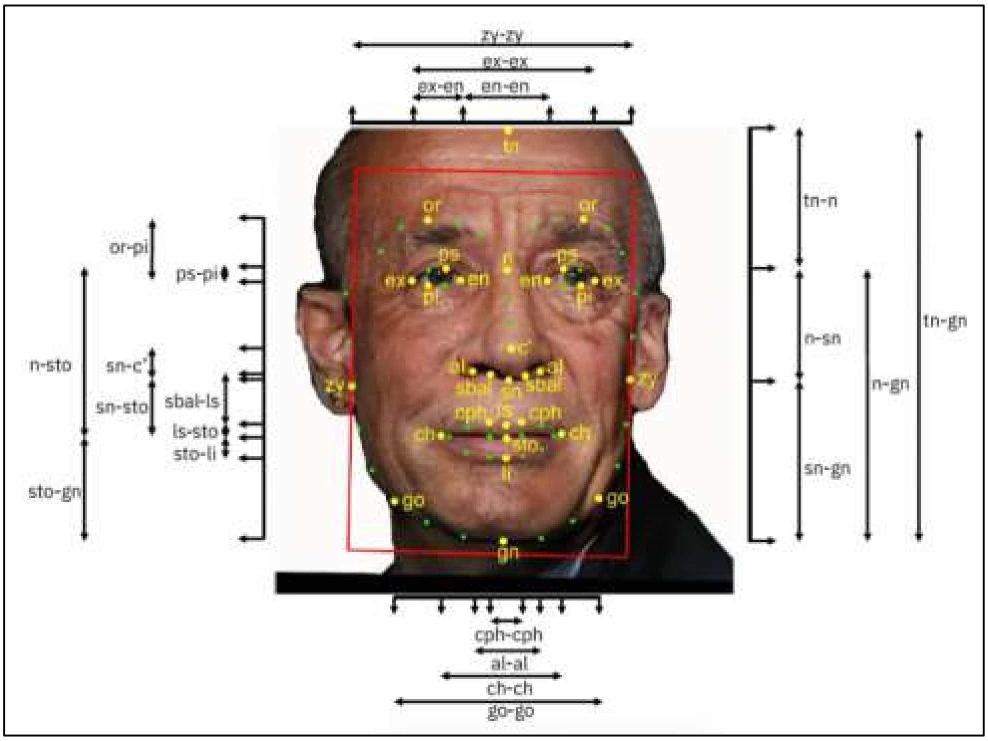

Figure 4 - Diversity dataset in Faces by IBM

Because facial recognition algorithms require a huge number of images, manual photo collection methods cannot meet their needs, so researchers have begun to collect images en masse from sites such as Flickr, Facebook , YouTube , and others.

The article [19] details the problem of unauthorized extraction of millions of photographs from the Internet to train corresponding facial recognition algorithms. In January 2019, IBM released a collection of almost a million photos from the Flickr platform, mentioning the intention to reduce objective errors using a diverse training dataset, which in turn caused a serious resonance among photographers whose works were included in this collection without notification of both the author and the model. In particular, concerns were expressed that the collected data could be used to restrict fundamental rights and privacy, as well as repressive and discriminatory policies.

Such random photos also become elements of the information field associated with certain objects, for example, with people whose faces can be identified in the photograph and then their identity can be established in more detail.

It is impossible to take such data into account in personal control, unlike the information that a person independently places in the public domain. In particular, people who closely work with the public remotely with the participation of cameras and microphones (even as amateurs) actually provide in the public domain recordings of their voice, their facial expressions, the vocabulary used and much more, which can be used for various purposes. For example, by telephone scammers for malicious machinations, who are able to "steal a voice" even with the help of a telephone conversation.

Users are already developing ideas to protect against such thefts, for example, developing voice avatars to protect subscribers when receiving calls from unknown persons [19].

When it comes to large commercial neural network image processing systems that work with government and security agencies, the situation is also not so clear-cut. In a groundbreaking 2018 study that had a significant impact on AI research, Joy Buolamwini and Timnit Gebru [20] were the first to find that all popular facial recognition systems were most accurate at identifying light-skinned men (2.4% error rate) and most likely to fail when recognizing dark-skinned women (61% error rate). Possible reasons for this phenomenon include the lower number of dark-skinned women in the databases, the predominantly white male composition of the developers of such systems, and the poor performance of camera sensors in recognizing details in dark shades. This is exacerbated by the fact that some commercial companies approach the development of neural network algorithms from a “black box” perspective, where they receive results and compare them with what they would like to receive without examining the essence of the internal processes.

Despite the identified problems, these systems continue to be widely used in various fields, including law enforcement agencies in Russia and China. Research confirms that members of racial minorities in these countries are at higher risk of being falsely identified as criminals. This trend is due to the fact that the system's algorithms are more likely to match faces whose features are similar to those of the suspect. A famous example of such a false match occurred in 2023, when a hydrologist was arrested in a 20-year-old murder case based on artificial intelligence (AI) data. According to news articles [21], the AI program determined that the detainee's photo was 55% similar to the image of a suspect in the 2003 murders. The case was closed only a year later.

This makes it clear that regulation is needed in the field of artificial intelligence (AI). Regulation should force AI developers to follow common standards so that they do not skimp on safety. Although regulations do not create technical solutions by themselves, they can still provide a powerful incentive to develop and implement them. Companies will develop safety measures more intensively if they cannot sell their products without them, especially if other companies are subject to the same standards. Some companies might regulate themselves, but government regulation helps prevent less careful competitors from skimping on safety. Regulation should be proactive, not reactive. It is often said that in aviation, regulations are “written in blood” - but here they need to be developed before a disaster, not after. They should be designed to give a competitive advantage to companies with better safety standards, not to companies with more resources and better lawyers. Regulators should be recruited independently, not from a single source of experts (e.g. large companies), so that they can focus on their mission for the common good without external influence [22, 23].

To increase transparency and accountability of AI systems, companies should be required to provide data documentation that explains what data sources they use to train and deploy their models. Companies' decisions to use datasets that contain personal data or invasive content increase the already frantic pace of AI development and hinder accountability. Documentation should describe the motivation for the choice, design, collection process, purpose, and maintenance of each dataset. Public oversight of general-purpose AI systems is also becoming increasingly important given the risks that private companies will never adequately consider. Direct public oversight of such systems may be needed to ensure that they are adequately addressed.

The ideal scenario would be that AIs are developed, tested, and then deployed only when all their catastrophic risks are negligible and under control. Before work can begin on a new generation of AI systems, the previous generation must undergo years of testing, monitoring, and deployment into society.

5. Conclusion

The rapid development of diffusion models has made it possible to achieve great results in various fields, from media to medicine and manufacturing. The ability of these models to generate realistic images and texts opens up new opportunities, but also creates many problems. One of the main points under consideration is the critical role of data and its quality in the process of training models, which directly affects their performance and accuracy. The concept of "information field" is introduced and substantiated. The need to pay attention to the issues of confidentiality and security of training data is noted, since with the increasing influence of neural networks, the risks also increase, requiring solutions to problems associated with the possible collapse of models and manipulation of public opinion through digital copies and pseudo-personalities. It is important to continue to research and develop neural network technologies, while implementing the necessary measures for their safe use.

6. Acknowledgments

The computational work was carried out using the K-100 hybrid supercomputer installed at the Center for Collective Use of the Keldysh Institute of Applied Mathematics of the Russian Academy of Sciences.

Refernces

1. I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde -Farley, S. Ozair , A. Courville , Y. Bengio , Generative Adversarial Networks, 2014, https://doi.org/10.48550/arXiv.1406.2661
2. J. Ho, A. Jain, P. Abbeel, Denoising Diffusion Probabilistic Models, 2020, https://doi.org/10.48550/arXiv.2006.11239 (date appeals 03/29/2023)
3. C. Meng, Y. He, Y. Song, J. Song, J. Wu, J. Zhu, S. Ermon, SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations, 2022, https://doi.org/10.48550/arXiv.2108.01073
4. R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-Resolution Image Synthesis with Latent Diffusion Models, 2021, https://doi.org/10.48550/arXiv.2112.10752
5. Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A.A., 2017. Image -toimage translation with conditional adversarial networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1125–1134.
6. Koh, J.Y., Baldridge, J., Lee, H., and Yang, Y., 2021. Text- image generation grounded by fine-grained user attention. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 237–246.

7. Ramesh A., Pavlov M., Goh G., Gray S., Voss C., Radford A., Chen M, Sutskever I., 2021. Zero-Shot Text-to-Image Generation, https://doi.org /10.48550/arXiv.2102.12092 D

8. Radford A., Jong WK, Hallacy C., Ramesh A., Goh G., Agarwal S., Sastry G., Askell A., Mishkin P., Clark J., Krueger G., Sutskever I. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv preprint arXiv:2103.00020 [cs.CV]. https://doi.org/10.48550/arXiv.2103.

9. Manda Bharadwaj, Dhayarkar Shubham , Mitheran Sai, VK, Viekash , Muthuganapathy , Ramanathan . 2021. ' CADSketchNet ' - An Annotated Sketch dataset for 3D CAD Model Retrieval with Deep Neural Networks. Computers & Graphics. 99. 10.1016/j.cag.2021.07.001.

10. Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez AN, Kaiser L., Polosukhin I. 2017. Attention Is All You Need, arXiv preprint arXiv:1706.03762 https://doi.org /10.48550/arXiv.1706.03762

11. Conceptual Captions Dataset. URL: https://github.com/google-research-datasets/conceptual-captions (date appeals 08/27/2024)

12. YFCC 100 M, URL: https://paperswithcode.com/dataset/yfcc100m (accessed 27.08.2024)

13. Generate CAD from text prompts, URL: https://zoo.dev/text-to-cad (date appeals 08/27/2024)

14. Bondareva N.A. Graphic neural networks and image verification problems // Proceedings of the 33rd International Conference on Computer Graphics and Machine Vision GraphiCon 2023, V.A. Trapeznikov Institute of Control Sciences of the Russian Academy of Sciences, Moscow, Russia, September 19-21, 2023, pp. 317-327, DOI: 10.20948/ graphicon - 2023-317-327 https: // www. graphicon. ru / html /2023/ papers / paper _031. pdf

15. Lomakina Ya. "The World's First Dead Blogger " Turned Out to Be a Living Actress: What Was That Anyway, 2024, URL: https://journal.tinkoff.ru/dead-blogger/ (date accessed 08/27/2024)

16. Clone yourself, 2024, URL: https://www.delphi.ai/ (accessed 27.08.2024)

17. Shumailov I., Shumaylov Z., Zhao Y., Gal Y, Papernot N., Anderson R., 2023. The Curse of Recursion: Training on Generated Data Makes Models Forget. arXiv preprint arXiv:2305.17493 https://doi.org/10.48550/arXiv.2305.17493

18. Rozhkov R. Gradation of degradation: is generative artificial intelligence facing degeneration, 2023, URL: https://www.forbes.ru/tekhnologii/491359-gradacia-degradacii-ozidaet-li-generativnyj-iskusstvennyj-intellekt-vyrozdenie (date of access 27.08. .2024)

19. Solon O. Facial recognition's 'dirty little secret': Millions of online photos scraped without consent, 2019, URL: https://www.nbcnews.com/tech/internet/facial-recognition-s-dirty-little-secret-millions -online-photos-scraped-n981921 (date appeals 08/27/2024)

20. Message from a telegram channel, 2023, https :// t. me / sburyi / 182 (date of access 08/26/2024)

21. Buolamwini J., Gebru T., Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification // Proceedings of the 1st Conference on Fairness, Accountability and Transparency, PMLR 81:77-91, 2018

22. TASS: Moscow court releases scientist arrested with AI in 20-year-old case, 2023, URL: https://tass.ru/proisshestviya/19508893 (date accessed 27.08.2024)

23. Kharitonova Yu.S., Savina V.S., Pagnini F. Bias of Artificial Intelligence Algorithms: Issues of Ethics and Law // Bulletin of Perm University. Legal Sciences. 2021. No. 53. URL : https: // cyberleninka.ru/article/n/predvzyatost-algoritmov-iskusstvennogo-intellekta-voprosy-etiki-i-prava (date of access : 09/05/2024) .

24. Hendrycks D., Mazeika M., Woodside T. An Overview of Catastrophic AI Risks arXiv preprint arXiv:2306.12001 https://doi.org/10.48550/arXiv.2306.12001