Digitalization and Management Innovation II
R.J. Dwyer (Ed.)

doi:10.3233/FAIA230736

Big Data Tools, Deep Learning & 3D Objects in the Metaverse

Preethi Ananthachari[1] and Storm Schutte
AI & Big Data Dept, Endicott College of International Studies, Woosong University, South Korea

Abstract. Metaverse is a virtual space to explore, for those who are not physically available in context. We are in the early stages of the development of the metaverse and it promises a rapid phase shift as well as proliferation in every aspect of technology. The metaverse comprises of eight key technologies, such as Extended Reality, Customer (user) Interactivity, Artificial intelligence, Block Chain Technology, Computer Vision, Internet of Things, Robotics, Fog/ Edge/ Cloud Computing, additionally, 5G/ 6G (Future) Mobile networks that are relevant to the metaverse. Big data is a significant source that complements each of these techniques. In this paper, we will discuss how different big data techniques help with the growth of the metaverse. The comparison of 3rd Dimensional objects in terms of Spark architecture with HDFS is carried out. We utilized large-scale 3D data to implement and evaluate various machine learning and deep learning models, facilitating a thorough comparison of their performance. Both of the Models have their own advantages and drawbacks, Random Forest model stands out in terms of processing speed, though SimpleNet requires more time for the validation accuracy improvement.

Keywords. Big Data, Deep Learning, 3D Objects, Spark, Metaverse

1. Introduction

Metaverse – hype in recent years is a word mixture of the Greek term ‘beyond’ (Meta) with the universe, describing a virtual world [1]. Technological people view the current trend as a positive transformation from the digital era to the metaverse universe. Although some people are skeptical about how it will support big data in the future. Every day there are numerous amounts of data being produced due to the usage of mobile phones, laptops, CCTV sensor data, etc,. Security is another big challenge in the metaverse. Metaverse is in its earlier stages of development however technological innovators can foresee its establishment in the future in a wider range [2]. Big, huge data production is inevitable in the future as an individual uses the internet for both personal and business purposes. How the data will be collected and processed in the metaverse is a big challenge in the future. As we know the world population is approximately 8 billion currently, and the number of connected devices in the world will increase by 18% about

[1] Corresponding author, Preethi Ananthachari, AI & Big Data Department, Endicott College of International Studies, Woosong University, Daejeon 34606, South Korea; E-mail: preethi.ga@wsu.ac.kr.

14.4 billion in 2023, also it is expected that there will be 27 billion connected devices by 2025 [3]. The global metaverse market is anticipated to reach 1.3 trillion by 2030 and its compound annual growth rate (CAGR) has been calculated as 44.5% up to 2030 by precedence research. The integration of mixed reality, blockchain, and machine/deep learning[4] will drive the digital world towards the metaverse. The pandemic of 2019 and 2020 significantly influenced the shift towards remote working practices. Even those unfamiliar with digital technology found themselves compelled to navigate it. A large number of IT companies adapted to this change by allowing their employees to work from home. As a result, people are actively equipping themselves to succeed in a world that is becoming more remote and increasingly virtual. Several platforms that cater to this new metaverse reality include Decentraland, Cryptovoxels, and Roblox, among others. Decentraland, for instance, is a virtual platform that enables users to purchase land using cryptocurrency. On the other hand, Roblox, widely recognized as the most frequently used metaverse platform, revolves around online games. Numerous companies have initiated experimental ventures with metaverse applications. A notable development in this field is the approval by the Food and Drug Administration (FDA) of a technology called Medivis. This Augmented Reality (AR) surgical system seamlessly integrates with a hospital's digital imaging system, allowing surgeons to rapidly synchronize their operations. [5]

Communication Transformation: Effective communication must be ensured, with a focus on seamless message transmission that is free from complexities. However, user interactions aren't limited to sending messages alone. With the integration of Virtual Reality (VR) and Augmented Reality (AR) technologies, alongside elements of personalization, voice interactions, and analog signals, a comprehensive brand experience can be crafted.

Consumer deep engagement: Virtual Reality (VR) profoundly immerses consumers in its innovative technology. From small businesses to large corporations, marketers are not only ready but also eager to integrate VR into their strategies. This application has the potential to engage customers more effectively, encouraging them to spend more time exploring products. As a result, this could significantly boost business growth. VR leans more towards visualizing stories rather than just telling them, aiding marketers in achieving their objectives more efficiently. Furthermore, VR cultivates enduring emotional engagement in consumers, thereby simplifying the seller's business operations.

Individual elements in metaverse: An individual is considered to be a digital avatar in 3D space. The currency used is a digital cryptocurrency which is based on blockchain technology. The environment is a virtual environment in which one can interact with others, do business etc,. One can create an account and invest in digital assets using cryptocurrency. Traveling to different locations is termed as teleportation in the metaverse. One can teleport to the next room or teleport to a nearby city, which is easy in the metaverse. The physical objects of the real world are represented as Digital twins in the metaverse. A Digital twin is the virtual form of a real physical object. When the avatars can visit different virtual worlds and exchange data in a seamless way, metaverse is interoperable [6,7].

2. Impacts and Growth of Big Data and Deep Learning

The global big data analytics market size is valued at $307.52 billion in 2023 and is estimated to grow to $745.15 billion by 2030 exhibiting a CAGR of 13.5%. Big Data analytics inspects structured and unstructured data to figure out and provide insights based on correlation, identify hidden patterns, track varying market trends, and more. Major public and private sectors intend to make use of analytical tools to obtain customer insights by developing strategic business analyses. [8] Past historical and real-time data can be analyzed together to evaluate the growing preferences of consumers or corporate customers, facilitating businesses to be more responsive to consumer requests and needs. Big data already made an establishment in healthcare by medical researchers to identify disease signs and risk factors and secures classified data for doctors to help diagnose illnesses and medical conditions in patients. Moreover, integrating data from digital health records, social media platforms, online resources, and other sources, equips healthcare organizations and government agencies with the most recent and relevant information regarding potential infectious disease threats or epidemics. Banking and finance sectors use big data systems for risk management and real-time analysis of market data. Business Product manufacturers and transportation companies rely on big data to manage, handle their supply chains and enhance delivery routes [9].

Chatbots are already making use of Deep learning models. Also, as it persists to mature, deep learning is expected to be employed in diverse businesses to improve customer experience as well as to get better customer satisfaction. AI Machines (models) are being taught the grammar and style of a set of text and are then using this model to automatically generate an entirely new narrative of text matching with the proper spelling, grammar as well as style of the original text. Tumor/Cancer researchers were already in progress, applying deep learning models into their practice as a way to automatically detect cancer cells. Deep learning models revolutionized the medical field using image processing by classifying the tumor region easily from MRI (magnetic resonance imaging) scans. It can easily identify millions of MRI scans in a very short span of time. [10] In the case of ChatGPT, deep learning is used to train the model's transformer architecture, which is a type of neural network that has been successful in various NLP tasks. The transformer architecture enables ChatGPT to understand and generate text in a way that is coherent and natural-sounding [11].

3. Related Work

The authors of [12] conducted both a primary survey and secondary research, which included expert articles from the Federal Trade Commission (FTC), Advertising Standard Council of India (ASCI), Meta Director, CXOs, Global Brand Consultants, Data Lawyers, Technology Leaders, and Senior Media Personnel. Their investigation revealed that the Metaverse is anticipated to bring about transformation in three main aspects – Consumer Experience, Resource Allocation, and Data Strategy.

The New York Times (NYT) in 2016, through its joint venture with Google, dispatched Google Cardboards to their online-only subscribers. Google Cardboard, a disposable VR headset, permitted the publication's digital-only subscribers to watch a VR film entitled 'Seeking Pluto's Frigid Heart', which allowed the users an intimate look at the dwarf planet. The NY Times also used Google Cardboard to tell the visual, deep/visceral story of the war's effects on children through an immersive documentary film entitled 'The Displaced', which won the Entertainment Grand Prix at the 2016 Cannes Lions Festival [13].

An attention-based approach utilizing Context R-CNN has been developed to enhance the performance of object detection on a specific frame. This strategy involves indexing into a long-term memory bank, established on a per-camera basis, and aggregating contextual features from other frames [14]. Concurrently, it's crucial to note that VR headsets collect sensitive biometric data. This data encompasses a broad range of intimate personal details, such as body posture, eye gaze, viewed images, pupil dilation, unfocused areas, gestures, touch sensations, interactions, spoken words, and even subtle changes in skin color or blushing [15].

In [16], authors built a segmentation model for image segmentation in which it transfers zero-shot to a new image distribution. In [17], authors proposed a hybrid model based on CNN, which inputs the features extracted by CNN into RF for classification. The author also claimed that the Random forest model has faster training and high classification accuracy. Authors in [18], used an ample number of fully connected networks accounting for the variance caused by random weight initialization.

4. Big Data

Big Data processes massive amounts of data. It collects the raw data, filters the unwanted part by focusing on the region of interest, and secures the intelligent data. By further processing, it identifies hidden patterns in big data, which is helpful in many fields. Mary Teren, a high school science teacher in Cobb County, Ga., revealed that in 30 minutes, a VR device collects millions of data points from their school students [15]. By this, one can imagine the amount of data generated in the metaverse universe. A High-end data processing tool will be required. Big data analytics is described as segregating, classifying, or grouping the varieties of data based on their key points. So, big data processing tools and techniques are indispensable. These tools will pave the way for the metaverse transformation. Metaversetroop suggested that a bandwidth of 1500 TPS (Transactions per second), 25 Mbps for live streaming, and 2.35GB for content resolution is advisable for the metaverse [19]. As the number of avatars grows, the resource usage increases automatically as well as the above-given statistics increase exponentially, which results in big data. Beyond 5G, a peer-to-peer (P2P) communication network is suggested for the metaverse considering the interaction of avatars. In P2P, resource utilization is entirely dependent on the number of participants and also it's a decentralized network, this would be recommended for metaverse communication.

4.1 Data Management

Hadoop Distributed File System (HDFS) is the core of Hadoop systems. HDFS manages the data in blocks that are spread over different data nodes clustered together and managed by a name node. One of the biggest challenges is fault tolerance, in which if a data node corrupts, the Hadoop system manages the data loss. The biggest Hadoop cluster consists of 4000 nodes and has a total capacity of 14 petabytes each. In order to keep the rate of metadata operations high, HDFS keeps the whole namespace in RAM. The namenode persistently stores the namespace image and its modification log (the journal) in external memory such as a local or a remote hard drive. If the namenode fails, its latest state can be restored by reading the namespace image and replaying the journal [20]. HDFS follows write once, read many policies. Multiple writes are not supported. HDFS follows a block mechanism for data storage and the size of the block is usually 128MB. By default it is 64MB. The fixed size makes it easier to handle the data.

DFS has two types of nodes in its cluster, namenode as well as datanode. It follows master-slave architecture in which the namenode is the master and manages all the datanodes. Datanodes contain blocks. The block information is maintained by the namespace in the namenode. Periodically this information will be updated when the datanode updates. Since the namenode maintains everything, if it fails, the whole system goes down. Namenode manages thousands of datanodes. Namenode requires about 150 bytes for managing a single datanode block. The number of datanodes handled by a single namenode, depends on the RAM size of the namenode. For security, a copy of the block file system can be maintained in a remote network file system (NFS). HDFS Federation allows more than one namenode. The namenodes will not communicate with each other. To overcome the namenode failure, in the High Availability concept, a pair of namenodes is maintained. One will be active and another will be on standby. In case of active namenode failure, the standby namenode takes charge. The transition from the active namenode to the standby namenode is said to be failover control. STONITH (Shoot the other node in the head) command will be carried out to stop the failed namenode from causing any damage by accessing other file systems. This is known as fencing.

4.2 Spark

Even Though Hadoop could run parallel jobs with many clusters, the design was a batch processing system and the MapReduce took much time for processing, Apache Spark [21] was founded as a variant for faster processing which uses the available RAM, and also has various libraries for processing streaming data. If big data is needed, the problem is that the data loss can be significant if a cluster node fails, so to prevent this we can use Apache Kafka (developed by LinkedIn in 2011 later donated to Apache Software Foundation) or Apache Flume, which will play the role of aggregation, storing log data, etc. It can be used as a means for streaming data which is important in the metaverse and similar applications.

Stream Processing: If the .obj files are being sent to your REST API in real time, Spark's Structured Streaming module can be used to process this data as it arrives. This allows the architecture to provide real-time predictions. ETL Operations: Spark can be used to extract, transform, and load (ETL) data. If the data comes in from multiple sources, in different formats, Spark can be used to clean, transform, and normalize the data before storing it in the Hadoop distributed system. Not to mention the role Spark could play in working with deep learning models [22].

5. 3D Object Detection

5.1 Dataset

In our study, we elected to employ .obj files as the medium for three-dimensional (3D) objects. This decision was informed by the ubiquity of the .obj file format, which allows for the storage of geometric data such as vertices, texture, polygons, meshes, and materials. Our dataset is primarily sourced from Objaverse [23], supplemented with our own data. Cumulatively, these datasets totaled approximately 11 terabytes. We meticulously curated the 3D objects, selecting those we envisaged could be useful in the metaverse or augmented reality. We partitioned the dataset into three categories: "Living Objects", "Small Man-Made Accessories", and "Larger Constructions". In total, the dataset consisted of 12,327 .obj files, equally distributed across these categories. For validation purposes, we set aside 20% of these files.

5.2 Data preparation

The majority of the data obtained is in .glbs format. For uniformity and to ensure compatibility with a common 3D object medium, we converted all files into the .obj format. Subsequently, we adjusted the 3D objects to comparable dimensions in terms of width, height, and length, effectively reducing their sizes, still keeping the data points for each of the vertices. Then meticulously cleaned the images, removing any superfluous mesh elements such as clouds, land, and materials unrelated to the object undergoing processing. Following this, 3D augmentations were implemented by to further enrich the dataset. The augmentation included object mutations, rotations, translations, and scaling. Moreover, random noise was injected into the object vertices to enhance the dataset's diversity. Finally, the faces of the 3D objects were converted into areas.

To fully leverage the capabilities of the selected architectures, we meticulously extracted the faces from the 3D objects and calculated the surface area of each individual face. This data was then cataloged into a CSV dataset, represented as a list for each object. This process was performed on an array of shapes, allowing us to accommodate a diverse range of calculations. Below fig:1 shows the representation of the 3D objects.

Figure 1. An example of four 3D objects used in our dataset. From the left, a man, woman and two houses.

Convolutional Neural Networks (CNNs): The rise of Convolutional Neural Networks (CNNs) has significantly impacted data science, particularly in the field of image classification [24]. Owing to its transformative potential, we strategically integrate CNNs with three unique architectures – the Dense Neural Network (DNN), AlexNet, and SimpleNet – to explore 3D object detection. Each of these architectural designs carries its own set of advantages and disadvantages, thereby providing a rich, multi-faceted understanding within the broader scope of our investigation.

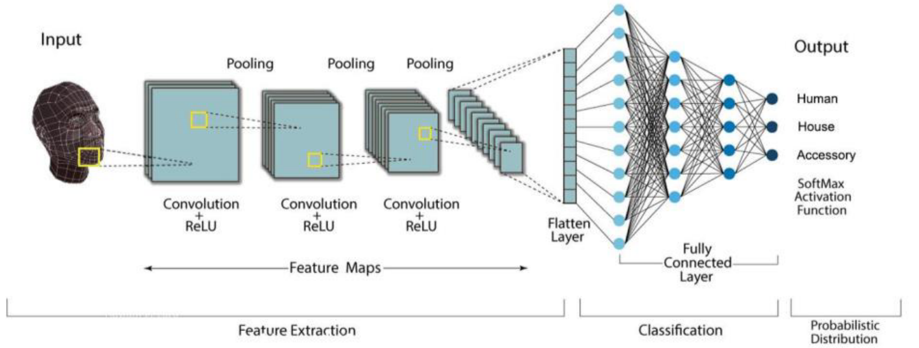

Figure 2. Convolutional Neural Networks Architect

Simpler Approaches with Alternative Machine Learning Models: In an effort to assess the viability of a more streamlined approach, six additional machine learning models were implemented. This assessment not only aimed to optimize data processing but also to mitigate computational costs. The selected models included the Random Forest, Support Vector Machine, Decision Tree, K-Nearest Neighbor (KNN), Gradient Boosting, AdaBoost, and a basic Neural Network. The application of these models provided valuable insights into whether a less complex, yet effective, strategy could potentially replace the need for more elaborate methods.

6. Results and Discussions

With the metaverse's rapid expansion, efficient categorization of user-generated and other 3D objects utilized within the metaverse or augmented realities are becoming paramount. Such categorization fuels diverse applications, from crafting personalized user experiences and effective Hadoop Distributed File System (HDFS) storage strategies to executing targeted advertising and enhancing gameplay narratives. Our research delves into the application of three eminent 3D convolutional neural networks and six machine learning models for processing and categorizing volumetric data.

Our primary challenge was to swiftly classify 3D objects to reduce computational costs and latency. To overcome this, we devised a strategy to extract each object's faces, calculate the area size of each face, and convert the .obj 3D image into easily processable numerical data, thus focusing on a specific aspect of the 3D object. Initially, we believed a deep learning model was essential for accurately predicting 3D objects. However, we discovered that by simplifying the data, conventional machine learning models were effectively able to manage basic classification tasks

A comparative evaluation of Dense Neural Net (Figure 5), SimpleNet (Figure 3), and AlexNet (Figure 4) revealed closely matched performance levels with regards to training accuracy and loss. While SimpleNet narrowly surpassed the others in accuracy, it registered a higher loss, hinting at a less-than-optimal fit. In terms of validation metrics, Dense Neural Net emerged as the potential leader with superior unseen data accuracy. Conversely, AlexNet, despite a lower validation accuracy, presented the lowest validation loss, indicating a slight advantage in generalizing to new data. On examining additional machine learning models, the Random Forest model stood out (Figure 6), offering unrivalled performance and computational efficiency. Other models, such as SVM, Decision Tree, K-Nearest Neighbors, Gradient Boosting, AdaBoost, and Neural Networks, frequently encountered challenges tied to high data dimensionality and computational constraints.

It's worth noting that our approach has limitations when it comes to extracting faces for which we can calculate the area. Some of the shapes embedded in the faces of the 3D objects are non-planar polygons which involves advanced geometric methods, requiring more advanced mathematical calculations to accurately measure the area within a 3D space. In conclusion, the Random Forest model emerged as the most effective method for processing big data, enabling swift and cost-effective classification of 3D objects.

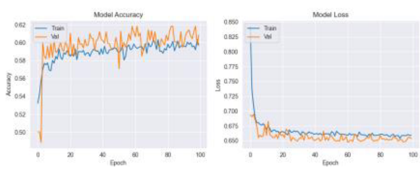

Figure 3. SimpleNet Model Accuracy and loss.

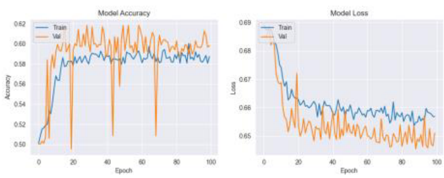

Figure 4. AlexNet Model Accuracy and loss.

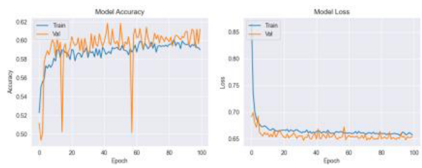

Figure 5. DNN Model Accuracy and loss.

Figure 6. Random Forest Confusion Matrix.

Table 1. Deep learning, machine learning statistics

Architect	CV Score	Tr. Acc.	Tr. Loss	F1 Score	Val. Acc.	Val. Loss
Dense Neuro Net.	n/a	0.590	0.657	n/a	0.612	0.654
SimpleNet	n/a	0.597	0.659	n/a	0.608	0.653
AlexNet	n/a	0.587	0.656	n/a	0.598	0.651
Random Forest	0.997	0.999	n/a	0.999	n/a	n/a
SVM	0.730	0.728	n/a	0.712	n/a	n/a
Decision Tree	0.984	0.988	n/a	0.988	n/a	n/a
KNN	0.981	0.989	n/a	0.989	n/a	n/a
AdaBoost	0.888	0.881	n/a	0.880	n/a	n/a
Gradient Boosting	0.979	0.979	n/a	0.979	n/a	n/a
Neural Network	0.884	0.890	n/a	0.888	n/a	n/a

7. Conclusion

This paper presented a method for preparing the metaverse to efficiently handle big data, with a specific emphasis on 3D images. In summary, our findings highlight the pivotal roles of HDFS as a distributed storage system for mapping collection data, and Spark for extracting faces from 3D objects and processing each face in preparation for prediction. Our research demonstrates that by implementing a strategy focused on 3D object faces, we can significantly reduce the data's size, enabling rapid processing through the

Random Forest model or other models, depending on the specific application. Our analysis revealed that machine learning models can achieve a good enough accuracy compared to their deep learning counterparts, even though each type of model has its own challenges with respect to loss and stability. The study underscores the importance of data simplification to prevent overfitting. Future research could explore the performance of different deep learning models using a broader range of 3D objects to optimize outcomes.

8. Future Work

In this paper, we have explored the potential these machine learning and big data tools provide in the rapidly evolving landscape where big data intersects with digital media in 3D spaces. Given the intriguing potential presented not only by digital media but also by three-dimensional biological imagery in the medical field, we are eager to explore in depth the integration of individual vertex identification together with the surface area measurement of polygons. By doing so, we anticipate a more nuanced outcome since the distinctiveness of each object would be exponentially magnified.

Furthermore, we recognize the need to introduce higher-dimensional geometric mathematics to accurately calculate the areas of non-planar polygons. The fusion of this mathematical approach with the aforementioned strategies should enhance precision, thereby requiring less computational power and significantly improving processing speed.

References

[1] Unpacking Meta: Where Did the Word Metaverse Come From? (2022, March 22nd) https://www.xrtoday.com/virtual-reality/unpacking-meta-where-did-the-word-metaverse-come-from/
[2] Olshannikova, E., Ometov, A., Koucheryavy, Y., & Olsson, T. (2015, October 1). Visualizing Big Data with augmented and virtual reality: challenges and research agenda - Journal of Big Data. SpringerOpen. https://doi.org/10.1186/s40537-015-0031-2
[3] How Many IoT Devices Are There in 2023? (n.d.). Techjury. https://techjury.net/blog/how-many-iot-devices-are-there/
[4] Yann LeCun1, Yoshua Bengio & Geoffrey Hinton, Deep learning (2015, May 28th) https://www.cs.toronto.edu/~hinton/absps/NatureDeepReview.pdf.
[5] Dmytro Spilka, Solvid "How Big Data will redefine our identity in the metaverse", https://venturebeat.com, 2022.
[6] Making the metaverse: hardware hurdles. (2021, Nov 25th). Making the Metaverse: Hardware Hurdles. https://www.electronicspecifier.com/products/artificial-intelligence/making-the-metaverse-hardware-hurdles
[7] What is the Metaverse? An Explanation and In-Depth Guide. (2023,.May 24th.). WhatIs.com. https://www.techtarget.com/whatis/feature/The-metaverse-explained-Everything-you-need-to-know
[8] January 2023, Big Data Analytics Market, Fortune business insight. https://www.fortunebusinessinsights.com/big-data-analytics-market-106179
[9] Bridget Botelho, Stephen J Bigelow, (January 2022), Big data, TechTarget. https://www.techtarget.com/searchdatamanagement/definition/big-data.
[10] Ed burns, Kate brush (March 2021), Deep Learning, TechTarget https://www.techtarget.com/searchenterpriseai/definition/deep-learning-deep-neural-network
[11] Alexandru Hutanu, (February 2023), Tech Trends, https://www.pentalog.com/blog/tech-trends/chatgpt-fundamentals/
[12] Krishnamurthy, R., Chawla, V., Venkatramani, A., & Jayan, G. (2022, December 18). Transforming Your Brand Using the Metaverse: Eight Strategic Elements to Plan For. Transforming Your Brand Using the Metaverse: Eight Strategic Elements to Plan for | California Management Review.

https://cmr.berkeley.edu/2022/12/transforming-your-brand-using-the-metaverse-eight-strategic-elements-to-plan-for/

[13] Mileva, G. (2022, April 13). Top 15 VR Marketing Examples for 2023. Influencer Marketing Hub. https://influencermarketinghub.com/vr-marketing-examples/

[14] Sara Beery , Guanhang Wu , Vivek Rathod , Ronny Votel , Jonathan Huang, "Context R-CNN: Long Term Temporal Context for Per-Camera Object Detection", 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020, DOI:10.1109/CVPR42600.2020.01309

[15] VR Devices Collect "Intimate" Data, Lack Privacy Protections. Should Schools Invest? (2022, November 18). Education Week. https://www.edweek.org/technology/vr-devices-collect-intimate-data-lack-privacy-protections-should-schools-invest/2022/11

[16] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W. Y., Dollár, P., & Girshick, R. (2023, April 5). Segment Anything. arXiv.org. https://arxiv.org/abs/2304.02643v1

[17] Erhui Xi, "Image Classification and Recognition Based on Deep Learning and Random Forest Algorithm", Wireless Communications and Mobile Computing, Volume 2022 | https://doi.org/10.1155/2022/2013181

[18] Leonardo F. S. Scabini, Odemir M. Bruno, "Structure and Performance of Fully Connected Neural Networks: Emerging Complex Network Properties", Physica A: Statistical Mechanics and its Applications, Volume 615, April 2023, 128585,ISSN 0378-4371, https://doi.org/10.1016/j.physa.2023.128585.

[19] Z. (2022, February 19). 5 Key requirements of metaverse. https://metaversetroop.com/requirements-of-metaverse/

[20] Konstantin V. Shvachko, "HDFS Scalability: the limits to growth", ;login: issue: ,April 2010, Volume 35, Number 2

[21] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, Ion Stoica, "Spark: cluster computing with working sets", HotCloud'10: Proceedings of the 2nd USENIX conference on Hot topics in cloud computing, June 2010, pp : 1~10.

[22] Deep Learning on Spark is Getting Interesting | Dell Technologies Info Hub. (2020, August 3). Deep Learning on Spark Is Getting Interesting | Dell Technologies Info Hub. https://infohub.delltechnologies.com/p/deep-learning-on-spark-is-getting-interesting/

[23] Deitke, Matt & Schwenk, Dustin & Salvador, Jordi & Weihs, Luca & Michel, Oscar & VanderBilt, Eli & Schmidt, Ludwig & Ehsani, Kiana & Kembhavi, Aniruddha & Farhadi, Ali. (2022). Objaverse: A Universe of Annotated 3D Objects.

[24] Rikiya Yamashita, Mizuho Nishio, Richard Kinh Gian Do & Kaori Togashi, "Convolutional neural networks: an overview and application in radiology" (2018, June 22nd) https://insightsimaging.springeropen.com/articles/10.1007/s13244-018-0639-9